

# EXHIBIT A

**UNITED STATES DISTRICT COURT  
WESTERN DISTRICT OF TEXAS  
AUSTIN DIVISION**

Intellectual Ventures I LLC and  
Intellectual Ventures II LLC,

Plaintiffs

v.

VMware, Inc.,

Defendant.

Civil Action No. 1:19-cv-01075-ADA

**PLAINTIFFS' CLAIM CONSTRUCTION BRIEF**

## TABLE OF CONTENTS

<b>I.</b>	<b>Introduction</b> .....	1
<b>II.</b>	<b>Legal Standards</b> .....	1
	<b>A. Claim Construction Generally</b> .....	1
	<b>B. Means-Plus-Function Claim Construction Analysis</b> .....	2
	<b>C. Indefiniteness</b> .....	2
<b>III.</b>	<b>The Disputed Claim Terms</b> .....	3
	<b>A. The '752 Patent</b> .....	3
	<b>B. The '686 and '726 Patents</b> .....	8
	1. “modify a resource allocation” / “modified resource allocation” / “modify[ing] [the] computer resources allocated to a virtual server” ('686 patent Claims 5, 6, 7) .....	9
	2. “modify[ing] a resource allocation for the virtual server” / “modifying [the] computer resources allocated to a virtual server” ('726 patent Claims 1, 4, 5 and 8) .....	9
	3. “resource denials” ('726 patent Claims 1, 4, 5 and 8) .....	10
	4. “resource unavailable messages” / “denied requests to modify a resource allocation” / “resource unavailable messages resulting from denied requests to modify a resource allocation” ('686 patent Claims 5, 6, 7) .....	11
	5. “virtual server” ('686 patent Claims 5, 6, 7); ('726 patent Claims 1-11) .....	12
	6. “quality of service guarantee” ('726 patent Claims 1 and 4) .....	15
	7. “determining that a second physical host can accommodate the requested modified resource allocation” ('686 patent Claims 5, 6 and 7) .....	16
	8. “determination that a virtual server is overloaded” ('686 patent Claims 5, 6 and 7) .	18
	9. “virtual server overload signal” ('726 patent Claims 1, 4, 5 and 8) .....	20
	10. Alleged Means Plus Function Elements ('686 patent Claim 7) ('726 patent Claims 1, 3, 5 and 7) .....	21
	<b>C. The '818 Patent</b> .....	29
	1. “virtual [network/storage network] interface layer of an application server” (Claims 1, 17, 32, 42) .....	29

2. “hierarchical token bucket resource allocation” / “token(s)” (Claims; 1, 17, 30, 32, 33, 37-42).....	31
3. “enforcing . . .” / “receiv[e/ing] . . .” / “classify[ing] . . .” / “compar[e/ing] . . .” / “forward[ing] . . .” / “buffer[ing] . . .” (Claims; 1, 17, 30, 32, 37-39, 42).....	33
4. “maintain[ing] a connection over a network fabric” (Claims 1, 17, 30, 32, 42).....	34
5. Alleged Means Plus Function Elements (Claim 17) (individually set forth in Ex. C)	36
<b>D. The ‘051 Patent</b> .....	38
1. “customer forwarding [table(s)/information]” (claims: 1, 3).....	38
2. “virtual server(s)” (claims; 1, 3, 6) .....	41
3. “physical interface(s)” (claims; 1, 3) .....	42
4. “storing . . .” / “determining . . .” / “using . . .” (claims; 1, 3) .....	43
<b>IV. Conclusion</b> .....	47

## TABLE OF AUTHORITIES

### Cases

<i>Acumed LLC v. Stryker Corp.</i> , 483 F.3d 800 (Fed. Cir. 2007).....	45
<i>Advanced Mktg. Sys., LLC v. CVS Pharmacy, Inc.</i> , 2016 WL 1741396 (E.D. Tex. May 3, 2016) .....	22
<i>Aloft Media, LLC v. Adobe Sys., Inc.</i> , 570 F. Supp. 2d 887 (E.D. Tex. 2008) .....	24, 26, 27
<i>Apple Inc. v. Motorola, Inc.</i> , 757 F.3d 1286 (Fed. Cir. 2014).....	37
<i>Bell Commc'ns. Research, Inc. v. Vitalink Commc'ns. Corp.</i> , 55 F.3d 615 (Fed. Cir. 1995).....	40
<i>Catalina Mktg. Int'l, Inc. v. Coolsavings.com, Inc.</i> , 289 F.3d 801 (Fed. Cir. 2002).....	33
<i>Collaborative Agreements, LLC v. Adobe Sys. Inc.</i> , 2015 WL 2250391 (W.D. Tex. May 12, 2015).....	23
<i>Comark Commc'ns, Inc. v. Harris Corp.</i> , 156 F.3d 1182 (Fed. Cir. 1998).....	9, 10, 14
<i>Cont'l Circuits LLC v. Intel Corp.</i> , 915 F.3d 788 (Fed. Cir. 2019).....	31
<i>CSB-Sys. Int'l, Inc. v. SAP America, Inc.</i> , 2011 WL 3240838 (E.D. Pa. July 28, 2011).....	23
<i>Dealertrack, Inc. v. Huber</i> , 674 F.3d 1315 (Fed. Cir. 2012).....	10
<i>Eolas Techs., Inc. v. Adobe Sys., Inc.</i> , 810 F. Supp. 2d 795 (E.D. Tex. 2011) .....	24, 26, 27
<i>Epistar Corp. v. Int'l Trade Comm'n</i> , 566 F.3d 1321 (Fed. Cir. 2009).....	35, 46
<i>Fisher-Rosemount Sys. v. ABB Ltd.</i> , 2019 WL 6830806 (S.D. Tex. Dec. 12, 2019) .....	21, 22
<i>Intellectual Ventures II LLC v. Bitco Gen. Ins. Corp.</i> , 2016 WL 125594 (E.D. Tex. Jan. 11, 2016) .....	24
<i>Inventio AG v. ThyssenKrupp Elevator Am. Corp.</i> , 649 F.3d 1350 (Fed. Cir. 2011).....	36
<i>Liebel-Flarsheim Co. v. Medrad, Inc.</i> , 358 F.3d 898 (Fed. Cir. 2004).....	19
<i>Markman v. Westview Instruments, Inc.</i> , 52 F.3d 967 (Fed. Cir. 1995).....	7
<i>Masco Corp. v. United States</i> , 303 F.3d 1316 (Fed. Cir. 2002).....	2
<i>Maurice Mitchell Innovations, L.P. v. Intel Corp.</i> , 2006 WL 1751779 (E.D. Tex. June 21, 2006).....	3, 4
<i>Micro Chem., Inc. v. Great Plains Chem. Co.</i> , 194 F.3d 1250 (Fed. Cir. 1999).....	7

<i>Nautilus, Inc. v. Biosig Instruments, Inc.</i> , 572 U.S. 898 (2014).....	3
<i>Netfuel, Inc. v. F5 Networks, Inc.</i> , 2017 WL 2834538 (N.D. Ill. June 29, 2017) .....	22
<i>O2 Micro Int’l Ltd. v. Beyond Innovation Tech. Co.</i> , 521 F.3d 1351 (Fed. Cir. 2008).....	45
<i>Phillips v. AWH Corp.</i> , 415 F.3d 1303 (Fed. Cir. 2005).....	1, 2, 7
<i>Pitney Bowes, Inc. v. Hewlett-Packard Co.</i> , 182 F.3d 1298 (Fed. Cir. 1999).....	10
<i>Retractable Techs., Inc. v. Becton, Dickinson &amp; Co.</i> , 653 F.3d 1296 (Fed. Cir. 2011).....	35, 46
<i>RLIS, Inc. v. Allscripts Healthcare Solutions, Inc.</i> , 2013 WL 3772472 (S.D. Tex. July 16, 2013).....	24, 26, 27
<i>Sonix Tech. Co. v. Publ’ns Int’l, Ltd.</i> , 844 F.3d 1370 (Fed. Cir. 2017).....	3
<i>Sound View Innovations, LLC v. Facebook, Inc.</i> , 2017 WL 2221177 (D. Del. May 19, 2017).....	22
<i>Thorner v. Sony Computer Entm’t Am. LLC</i> , 669 F.3d 1362 (Fed. Cir. 2012).....	2
<i>Vitronics Corp. v. Conceptronic, Inc.</i> , 90 F.3d 1576 (Fed. Cir. 1996).....	1
<i>Watts v. XL Sys., Inc.</i> , 232 F.3d 877 (Fed. Cir. 2000).....	2, 36
<i>Williamson v. Citrix Online, LLC</i> , 792 F.3d 1339 (Fed. Cir. 2015).....	2, 21, 38
<i>Zeroclick, LLC v. Apple Inc.</i> , 891 F.3d 1003 (Fed. Cir. 2018).....	38

#### Other Authorities

IEEE Standard Dictionary of Electrical and Electronics Terms, p. 1182 (6th ed. 1996) .....	14
--	----

Per the Court’s Order Governing Proceedings – Patent Case, plaintiffs, Intellectual Ventures I LLC and Intellectual Ventures II LLC (together “IV”), respectfully submit this claim construction brief in support of their proposed constructions for the disputed terms.

## **I. Introduction**

IV alleges that VMware infringes five patents owned by IV: United States Patent Nos. 7,949,752 (“the ’752 patent”); RE 44,686 (“the ’686 patent”); RE 42,726 (“the ’726 patent”); RE 43,051 (“the ’051 patent”) and RE 44,818 (“the ’818 patent”).<sup>1</sup> IV’s patents teach systems and methods to virtualize data centers and enable the use of cloud computing to deliver computer services. VMware’s products and services deliver virtualization and cloud computing services to its customers.

## **II. Legal Standards**

### **A. Claim Construction Generally**

In resolving claim construction disputes, the Court considers three “intrinsic” sources: (1) the claims; (2) the patent specifications; and (3) the prosecution histories. *Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1582 (Fed. Cir. 1996). Courts may also rely on extrinsic evidence, such as dictionaries or treatises. *Phillips v. AWH Corp.*, 415 F.3d 1303, 1317 (Fed. Cir. 2005). A trial court is not required to follow any set formula to construe disputed claim terms “[n]or is the court barred from considering any particular sources or required to analyze sources in any specific sequence . . .” *Id.* at 1324.

“The claims of a patent define the invention to which the patentee is entitled the right to exclude.” *Id.* at 1312. “The words of a claim are generally given their ordinary and customary meaning as understood by a person of ordinary skill in the art when read in the context of the

---

<sup>1</sup> References in this brief to “Ex. \_\_” are to exhibits attached to the Declaration of Jonathan R. DeBlois (“DeBlois Decl.”).

specification and prosecution history.” *Thorner v. Sony Computer Entm’t Am. LLC*, 669 F.3d 1362, 1365 (Fed. Cir. 2012). *See also Phillips*, 415 F.3d at 1315 (“In some cases, the ordinary meaning of the claim language... may be readily apparent even to lay judges, and claim construction in such cases involves little more than the application of the widely accepted meaning of commonly understood words.”). There are only two exceptions to the general rule that claims are given their ordinary and customary meaning: “1) when a patentee sets out a definition and acts as his own lexicographer, or 2) when the patentee disavows the full scope of a claim term either in the specification or during prosecution.” *Thorner*, 669 F.3d at 1365. Extrinsic evidence may be consulted and used to construe claim terms only if such evidence does not contradict the intrinsic evidence. *Phillips*, 415 F.3d at 1322-23.

#### **B. Means-Plus-Function Claim Construction Analysis**

Section 112, Paragraph 6<sup>2</sup> provides that a structure may be claimed as a “means...for performing a specified function.” *Masco Corp. v. United States*, 303 F.3d 1316, 1326 (Fed. Cir. 2002). If the claim language does not include the word “means” then there is a rebuttable presumption that § 112, ¶ 6 does not apply. *Williamson v. Citrix Online, LLC*, 792 F.3d 1339, 1348 (Fed. Cir. 2015). To overcome this presumption, it must be shown that the term at issue fails to “recite sufficiently definite structure” or “function without reciting sufficient structure for performing that function.” *Watts v. XL Sys., Inc.*, 232 F.3d 877, 880 (Fed. Cir. 2000).

#### **C. Indefiniteness**

The Patent Act requires claims to particularly point out and distinctly claim the subject matter regarded as the inventions. 35 U.S.C. § 112, ¶ 2. To satisfy this requirement, the claim must be read in light of the intrinsic evidence to determine whether it informs one of skill in the

---

<sup>2</sup> The Pre-AIA version of § 112, ¶ 6 applies to the claims in this case.



art at the time of the invention “about the scope of the invention with reasonable certainty.”

*Nautilus, Inc. v. Biosig Instruments, Inc.*, 572 U.S. 898, 910-11 (2014). To establish that a claim

is indefinite, a patent challenger must prove indefiniteness by clear and convincing evidence.

*Sonix Tech. Co. v. Publ’ns Int’l, Ltd.*, 844 F.3d 1370, 1377 (Fed. Cir. 2017).

### **III. The Disputed Claim Terms**

#### **A. The ’752 Patent**

The ’752 patent was issued on May 24, 2011 and claims priority to an application filed on October 23, 1998. The inventions described and claimed in the ’752 patent enable the delivery of highly reliable and customizable cloud and virtualization services to customers whose local computing platforms can be simple enough to support only a web browser.

The litigation history of the ’752 patent has particular relevance to the disputed terms now before the Court. Specifically, on July 8, 2015, IV filed a patent infringement suit against HCC Insurance Holdings, Inc. (Case No. 6:15-cv-660) (hereinafter “HCC case”) that included the ’752 patent. On August 26, 2016, Magistrate Judge Mitchell entered a Report and Recommendation construing certain terms of the ’752 patent. *See* HCC case, Dkt. 102 (hereinafter the “R&R”). Both HCC and IV filed objections to the recommended constructions, however, the case was resolved prior to a decision on the parties’ objections.

In the interest of conserving Party and Court resources, IV has adopted many of the constructions set forth in Judge Mitchell’s R&R and agrees with the reasoning supporting her conclusions. While the R&R is not determinative, it is considered persuasive. *See Maurice Mitchell Innovations, L.P. v. Intel Corp.*, 2006 WL 1751779, at \*3-4 (E.D. Tex. June 21, 2006). VMware, on the other hand, agreed to adopt only one of the prior court’s constructions, and in many cases proposes the very same constructions that Magistrate Judge Mitchell rejected.

1. “exhausted” (Claims 1, 9 and 24)

IV’s Proposed Construction	VMware’s Proposed Construction
“used up to the allotted or pre-determined amount”	“unavailable for reuse”

VMware requests the Court to construe this disputed term in a manner inconsistent with the intrinsic record and contrary to the prior recommended construction from Judge Mitchell’s R&R. *See Maurice Mitchell Innovations, L.P.*, 2006 WL 1751779, at \*4 (treating the prior court’s analysis as persuasive is consistent with “the premise that a uniform treatment of claim construction is desirable”).

Judge Mitchell’s logic stands true and should be followed. The intrinsic record is replete with evidence supporting IV’s proposed construction and, in contrast, nowhere unambiguously states that the disputed term is limited to being unavailable for reuse as VMware suggests. For example, the specification discloses that resource consumption by agent 22 is monitored and further consumption is halted when the amount of said resource held by agent 22 is exhausted. Ex. D at 52:31-34. The disclosure goes on to say that each agent has permission to consume up to a pre-authorized amount of each service resource when performing a task. *Id.* at 9:40-43. Further still the patent states that a service wrapper can be used to ensure that the agent 22 does not consume more than its allotted amount of any particular service resource as specified by a respective service permission. *See, e.g., id.* at 22:34-37, 22:47-50, 24:63-67, 25:1-6. One of ordinary skill in the art would understand that “exhausted” in the context of the ’752 patent means “used up to the allotted or pre-determined amount.”

IV’s construction also reflects the overall goal and novelty of the inventions embodied in the ’752 patent. The patent describes a customizable and virtualized solution that enables service provider customers to utilize the services made available by the provider in a way that avoids a “one size fits all” solution. In other words, it allows customers to pay for compute power, time

and access on an as needed basis. In this context were “exhausted” to mean “unavailable for reuse” as VMware proposes, the solution would fail at the most basic level. One customer using an amount of a service and service resource would permanently make that service and service resource unavailable to subsequent customers, even after he or she has released the resource and is no longer using it. Narrowing the term in this manner is not supported by the intrinsic record.

2. “consumed”/“consumption” (Claims 1, 9, 22 and 24)

IV’s Proposed Construction	VMware’s Proposed Construction
“used”	“used up”

IV’s proposal of construing the disputed term as “used” reflects the patentees’ intent as extensively disclosed in the specification. IV can find only a single instance where “used” is referenced alternatively as “used up.” *See* Ex. D at 8:21-23. In contrast, the patent uses “consumed” and “used” interchangeably at least half a dozen times. *See, e.g., id.* at 12:25-30 (“ . . which may also monitor the amount of each respective service resource **consumed** to ensure that no particular agent **uses** more than an amount authorized . . .”), 16:50-55 (“ . . . monitor the amount of respective service resources expended, **used**, or otherwise **consumed** by one or more agents”), 25:14-16 (“at step 818, service wrapper 26 identifies the amount of each service resource actually **consumed** or **used** to execute the instruction.”) (emphasis added).

The citation in a single sentence that references “consumed” as “used up” is insufficient to support VMware’s construction, particularly in light of the widespread use of “consumed” and “used” as interchangeable. Therefore, the Court should reject VMware’s unsupported position and construe the term “consumed” as “used.”

3. “service” (Claims 1, 3, 9, 14 and 24)

IV’s Proposed Construction	VMware’s Proposed Construction
“network functionality available to agents”	“an application that is used by an agent on behalf of a principal”

The first description of the disputed term appears in the Abstract and discloses that “an agent is operable to utilize a service within the network system.” This description of the service as functionality of the network system that is available to the claimed agents is consistent throughout the entire patent. In fact, there is only a single mention of one embodiment of the patented invention where it’s noted that services may comprise a software application available to a principal which may be used by an agent on its behalf. Ex. D at 10:17-20. Presumably it is from this lone citation which VMware attempts to read into “services” the proposed limitations noted above. Such a limitation, however, is not only contrary to the basic principles of claim construction but is overwhelmingly discredited by the remaining 35 pages of description and disclosure in the ’752 patent.

For instance, the specification describes the functionality of a service as not limited to a particular form, such as an application, but rather, various network functionality made available by the operator or “service provider” to the claimed agents in order to perform an operation or task. *See, e.g., id.* at 12:14-25 (“[i]n operation, one or more agents may be set up for each user who is a subscriber to the services offered by the operator/provider of the network system . . . [e]ach agent for a particular user performs one or more tasks on behalf of that user . . . [t]o perform these tasks, each agent utilizes one or more services, during which it may consume various respective service resources.”). Examples of the claimed “services” available via the network system to the subscribers include “an e-mail service, a voice mail service, a paging/facsimile service, an address book and calendar service, and a business news and stocks information service, respectively.” *Id.* at 13:21-28, Fig. 2. Notably, these are not described as “applications” but are said to be “sub-systems” of the network system coupled to fast Ethernet hubs, and data storage and processing hardware to support the aforementioned

services. *Id.* at 13:15-20. While IV does not disagree that the disclosed services could be compiled into discrete programs, it does dispute that the specification limits the claimed services in such a way.

A further confirmation of IV's position that the claimed "services" are not limited to applications, but in fact represent functionality of the networked system, can be found in the Microsoft Computer Dictionary, p. 475 (5th ed. 2002) [hereinafter Ex. O] (emphasis added)<sup>3</sup>, which defines "services" in the context of networking as "specialized, software-based **functionality** provided by network servers." (emphasis added). This definition is consistent with the disclosures in the specification cited above and further evidences that IV's construction should be adopted.

#### 4. Means Plus Function Elements (Claims 1, 3, 4, 6) (individually set forth in Ex. B)

As an initial matter, VMware attempts to impermissibly narrow the identified structure for each and every functional clause in Ex. B by qualifying the corresponding structure with the language "as described in" and citing cherry-picked portions of the specification thereafter. This is improper. Identifying the specific structure corresponding to the claimed function is all that is required. Any further attempt to limit the corresponding structure, for example, by limiting that structure to only a single embodiment, is not allowed. *See Micro Chem., Inc. v. Great Plains Chem. Co.*, 194 F.3d 1250, 1258 (Fed. Cir. 1999) ("Identification of corresponding structure may embrace more than the preferred embodiment.").

As one example, for the element "*means for mediating an interaction between the means for using the service and the service*," both parties identify "service wrapper 26" as

---

<sup>3</sup> Although extrinsic evidence is not dispositive, contemporaneous evidence of the understanding of one of skill in the art is permitted to further support the intrinsic record. *See Phillips*, 415 F.3d at 1317-18. *See also Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 980-81 (Fed. Cir. 1995).

corresponding structure. VMware, however, further limits that structure to “as described in 16:22-38.” This reads out several other instances in the specification where service wrapper 26 is describe as performing the claimed function. *See, e.g.*, Ex. D at 3:20-27, 17:43-46, 18:49-54, 25:1-24. Accordingly, the Court should disregard VMware’s narrowing pin cites.

Aside from the impermissible qualifying of corresponding structure discussed above, the parties only disagreement relates to clause 6 in Ex. B. IV and VMware are in agreement on the claimed function of clause 6, but the structure identified by VMware is too narrow and reads out broader structure clearly linked to performance of the claimed function. For example, VMware proposes that the structure is “monitor (50),” yet monitor 50 is only a sub-part of service wrapper 26—proposed by IV—that collectively performs the claimed function of “monitoring an amount of the service resource used by the network-based agent.” The specification, for instance, discloses that “service wrapper 26 identifies the amount of each service resource actually consumed or used to execute the instruction . . . [then] asks agent server to decrement the amount allotted to agent 22 by the amount actually used.” Ex. D at 25:14-18. So while IV agrees that monitor 50 plays a role in performing the claimed function (as it is a sub-component of service wrapper 26), it is clear from the above citation that service wrapper 26 as a whole is also involved. Therefore, the Court should adopt IV’s proposed structure corresponding to the function of clause 6 because it properly encompasses the entire scope of the corresponding structure.

#### **B. The ’686 and ’726 Patents<sup>4</sup>**

While the ’686 patent issued on December 31, 2013, and the ’726 patent issued on September 20, 2011, both claim priority to an application filed on May 11, 2000. The inventions

---

<sup>4</sup> Because the two patents share a specification, unless otherwise noted all citations in this Section are to the ’726 patent.

described and claimed in the '686 and '726 patents relate to systems and methods for transferring virtual machines from one physical server to another physical server to help avoid outages, and thereby enable increased service continuity for business and consumer customers.

1. “modify a resource allocation” / “modified resource allocation” / “modify[ing] [the] computer resources allocated to a virtual server” ('686 patent Claims 5, 6, 7)
2. “modify[ing] a resource allocation for the virtual server” / “modifying [the] computer resources allocated to a virtual server” ('726 patent Claims 1, 4, 5 and 8)

IV's Proposed Construction	VMWare's Proposed Construction
“modif[y/ied] set of functions and features of a physical host(s) used in implementing tasks for each virtual server” / “modify[ing] a set of the functions and features of a physical host(s) used in implementing tasks for each virtual machine”	“modif[y/ied] a quality of service guarantee” / “modify[ing] [the] computer resources allocated to a virtual server”

IV's proposal for construing the disputed terms should be adopted because, (a) it correctly identifies and incorporates the intrinsic definition of “resource/resource allocation,” and (b) it accurately reflects the knowledge of one of skill in the art. VMware's proposal, on the other hand, simply reads in “quality of service guarantee” while simultaneously reading out “resource allocation.” Not only is this not helpful, but it's contrary to black letter claim construction law, and therefore, should be rejected. *See Comark Commc'ns, Inc. v. Harris Corp.*, 156 F.3d 1182, 1186-87 (Fed. Cir. 1998).

The disputed terms at issue here appear in seven claims across both the '726 and '686 patents. Since “modify” is well known and used in its plain meaning, the construction of the “resource allocation” iterations are the real dispute. IV's proposal takes its support directly from the intrinsic record which clearly and unambiguously defines “resource” as “the set of functions and features the physical host machine uses in implementing tasks for each virtual server.” Ex. F at 4:28-30. By directly incorporating the specification's definition of the disputed terms into its

proposed construction, IV's proposal captures the true scope of the disputed terms while following the well-known claim construction principals of not reading extraneous limitations into the claims. *See Comark Commc 'ns, Inc.*, 156 F.3d at 1187.

Notably, VMware reads out "resource allocation" entirely and replaces it with "quality of service guarantee," a term found only in the preamble of two of the seven claims at issue. Even were VMware arguing that the preamble of those two claims is limiting (which it is not), the preamble is not necessary to breathe life into the claims, and thus cannot be limiting. *Pitney Bowes, Inc. v. Hewlett-Packard Co.*, 182 F.3d 1298, 1305 (Fed. Cir. 1999). Accordingly, the Court should reject VMware's transparent attempt to re-write the claims and adopt IV's proposed construction for these disputed terms.

3. "resource denials" ('726 patent Claims 1, 4, 5 and 8)

IV's Proposed Construction	VMWare's Proposed Construction
"an indication that a request by the virtual server cannot be immediately serviced"	"an indication that a request by the virtual server for additional resources is either implicitly or explicitly denied"

The Court should adopt IV's proposed construction of the disputed term because "resource denials" is broadly defined in the specification and VMware's proposed construction is merely one narrow example of a single preferred embodiment. More specifically, IV's proposed construction is a direct quotation from the portion of the specification where the invention as a whole, rather than a specific embodiment, is described. Ex. F at 2:55-56 ("a resource denial **may refer to any request by the virtual server that cannot be immediately serviced**") (emphasis added). Such disclosure clearly indicates that the use of the disputed term was not intended to be limited beyond this general definition. *See Dealertrack, Inc. v. Huber*, 674 F.3d 1315, 1321-22 (Fed. Cir. 2012).



Despite black letter claim construction law directing that a term’s construction should not be limited to specific exemplary embodiments, VMware nevertheless has chosen to propose a construction that appears in the specification as an illustrative embodiment. For instance, the portion of the specification that VMware is presumably using as support for its proposal is discussing Figure 2A and notes in prefatory language that it’s “an embodiment of the overall process . . .” and only “one embodiment.” Ex. F at 5:21-34. Further down in the specification within the discussion of yet another embodiment reflected in Figure 3, VMware’s proposed construction is again noted as “an embodiment of one process for determining whether an individual resource in a virtual server has reached its resource limit . . .” and “FIG. 3 shows four examples of resource denial signals.” *Id.* at 7:41-43, 7:51-55. Accordingly, because VMware is trying to read one preferred embodiment into the disputed term in an instance where it’s otherwise clear that the disputed term is not so limited, the Court should adopt IV’s proposal.

4. “resource unavailable messages” / “denied requests to modify a resource allocation” / “resource unavailable messages resulting from denied requests to modify a resource allocation” (’686 patent Claims 5, 6, 7)

IV’s Proposed Construction	VMWare’s Proposed Construction
“an indication that a request by the virtual server cannot be immediately serviced” / “a request by the virtual server that cannot be immediately serviced”	“messages that indicate that a request to modify a resource allocation has been denied;” <i>see</i> construction of “modify a resource allocation”

As an initial matter, IV has proposed “resource unavailable messages” and “denied requests to modify a resource allocation” for construction, while VMware proposes “resource unavailable messages resulting from denied requests to modify a resource allocation.” In an effort to avoid confusion as much as possible IV will treat its two proposed terms as distinct parts of the larger term proposed by VMware.

IV’s construction stays true to the specific sequence of the claim element as a whole. As written the claims require a resource unavailable message be the result of “denied requests to

modify a resource allocation.” *See, e.g.*, Ex. E at Claim 5(a). In the summary of the invention the patentee describes how such a determination is made, disclosing that resource denials are monitored and used to determine whether the virtual server is overloaded. *See id.* at 2:62-65. The patentee then gives a general definition of “resource denials.” *See id.* at 65-66 ([a] resource denial may refer to any request by the virtual server that cannot be immediately serviced”). Thus, the disputed terms should be read as the result of a resource denial, i.e., an indication that a request by the virtual server cannot be immediately serviced.

VMware’s proposal merely rearrange the claim language. This is not only unhelpful but also completely ignores the sequence of the claim language and disclosure reflected in the intrinsic record discussed above. Therefore, since IV’s proposed construction tracks the language and sequence in the claims as well as the teachings of the specification, it should be adopted by the Court.

5. “virtual server” (’686 patent Claims 5, 6, 7); (’726 patent Claims 1-11)

IV’s Proposed Construction	VMWare’s Proposed Construction
Plain and ordinary meaning; in the alternative, “a virtual machine that resides on a physical server and uses the physical server’s resources, but has the appearance of being a separate, dedicated machine”	“a virtual server capable of receiving a quality of service guarantee from a physical host”

The disputed term “virtual server” should be given its plain and ordinary meaning. The term’s use in the specification and throughout the claims is consistent with its customary use as would be evident to one of skill in the art at the time of the invention. For instance, the background section of the specification begins by stating that the field of the invention relates to “resource allocation for a virtual server . . . and dynamically modifying the resource allocation for a virtual server. . .”). Ex. F at 1:24-27. Taking this statement at face value, if the field of art in question is resource allocation among virtual servers, one of skill in that art would readily

know what a virtual server is. Furthermore, the remainder of the background section describes the technology at issue in well-known and commonly understood terms. *See, e.g.*, Ex. F at 1:43-46 (“an ISP would prefer to offer network services to multiple customers while keeping all of the server host computers within a central location of the ISP . . .”), 1:53-55 (“most customers will neither require nor be amenable to paying for the [use] of an entire host computer”), 2:3-5 (“when servicing the needs for multiple customers having different needs, it is desirable to provide a virtual server that is dynamic, not static, in its allocation of resources”). It is clear from these disclosures that the patentee is using the term “virtual server” as it was understood in the art at the time, and therefore, should be construed according to its plain and ordinary meaning.

The file history of U.S. Pat. No. 6,985,937 (the originally issued parent of the ’726 and ’686 patents) [hereinafter “’937 patent”] similarly makes clear that the term “virtual server” is being used in its customary manner. During the prosecution of the aforementioned patent, the patentee describes his invention’s novelty over prior art virtual servers. He states that “Yu mentions that classes can be mapped to virtual servers and that these virtual servers can then be mapped to physical servers . . . however, Yu does not suggest ‘a virtual server being allocated a portion of the resources’ of a physical host.” Ex. P at File History, Oct. 5, 2004 Office Action Response, p. 14. The patentee goes on to note that “[w]hile this implies the existence of a physical server to which multiple virtual servers are mapped, Yu does not disclose, teach, or suggest allocating to a virtual server a portion of the physical server’s resources.” *Id.* at 15. In making these statements, the patentee acknowledges that virtual servers are known in the art and that he is using the term consistent with its known meaning, and not some unconventional or heretofore unknown construct he is coining a “virtual server.”

VMware's proposal, on the other hand, should be rejected. In addition to violating the oft-cited principal that terms should normally be given their plain and ordinary meaning absent clear evidence to the contrary, VMware's proposal attempts to read in a narrowing limitation that is unsupported by the intrinsic evidence. For example, VMware's proposal uses the disputed term itself ("virtual server") and simply tacks on the limitation "capable of receiving a quality of service guarantee," a phrase found in the preamble of only two of the thirteen claims at issue. As an initial matter, there is no evidence that the preamble of these two claims is limiting, and in any case, even if it were, reading the limitation of one embodiment into the construction of the disputed term is improper. *See Comark Commc 'ns, Inc.*, 156 F.3d at 1187.

Should the Court be inclined to construe the disputed term as opposed to giving it its plain and ordinary meaning, then the term should be construed as "a virtual machine that resides on a physical server and uses the physical server's resources, but has the appearance of being a separate, dedicated machine." This is supported by the intrinsic record and contemporaneous extrinsic evidence as well. The patent describes the claimed virtual servers as being "hosted" by a physical server which "dedicate[s]" portions of its physical resources to servicing the "resident" virtual servers. Ex. F at 2:42-46. Furthermore, the specification states that a customer otherwise required to purchase a full server from an ISP could avoid doing so by instead purchasing a virtual server backed by a dedicated portion of a physical host server's resources. In other words, from the customer's point of view, the virtual server functions just as an entire physical host server would. *Id.* at 1:49-57, 2:3-17. Contemporaneous dictionaries buttress the intrinsic evidence in this regard. *See* Ex. O, p. 555 (defining virtual server as "a virtual machine that resides on an HTTP server but has the appearance to the user of being a separate HTTP server"); IEEE Standard Dictionary of Electrical and Electronics Terms, p. 1182 (6th ed. 1996)

(Ex. K) (defining a virtual machine as “a functional simulation of a computer and its associated devices”). Therefore, if the Court finds that the disputed term should be construed its plain and ordinary meaning then it should be construed as IV proposes.

6. “quality of service guarantee” (’726 patent Claims 1 and 4)

IV’s Proposed Construction	VMWare’s Proposed Construction
“a guaranteed resource allotment which can be dynamically increased/modified”	“information that specifies a guaranteed amount of an assigned resource, and that can be dynamically modified” <sup>5</sup>

The disputed term should not be limited to “information that specifies” or “an assigned resource,” as VMware suggests, because the intrinsic record is devoid of any such limitation. Rather, as IV proposes, the term simply refers to a guaranteed resource allotment capable of being dynamically modified. The patent describes the disputed term as requiring only two things, (a) to guarantee a resource allotment; and (b) that the resource allotment be capable of being dynamically adjusted. *See* Ex. F at 2:38-46 (“[t]he present invention **dynamically adjusts** the quality of service guarantees for virtual servers **based upon the resource demands** experienced by the virtual servers.”) (emphasis added). By inserting “information that specifies” into its proposal VMware is conflating two distinct concepts from the specification. Quality of service guarantee information is not the same thing as the disputed term. The specification discloses that **quality of service guarantee information** can be, for example, stored in a **quality of service parameter table** in each physical host. 4:47-49. The quality of service guarantee information, therefore, is merely a representation of the quality of service guarantee, not the quality of service guarantee itself. Therefore, VMware’s construction is simply incorrect.

---

<sup>5</sup> VMware substantially changed its proposed construction of this term just hours before the parties were to file their respective briefs, prejudicing IV’s preparation of its brief.

In addition, nowhere in the specification is a quality of service guarantee described as being limited to “an assigned” resource. By inserting the term “assigned” VMware is attempting to limit the disputed term in a way not consistent with the intrinsic record.

7. “determining that a second physical host can accommodate the requested modified resource allocation” (’686 patent Claims 5, 6 and 7)

IV’s Proposed Construction	VMWare’s Proposed Construction
plain and ordinary meaning; in the alternative, “determining that a second physical host can accommodate the request(s) by the virtual server that could not be serviced immediately”	indefinite; in the alternative, “determining that a second physical host can accommodate the denied requests to modify a resource allocation”

The disputed term should be given its plain and ordinary meaning in light of the patent’s specification, which when read—even cursorily—illustrates that the disputed term is not indefinite as VMware suggests. The intrinsic record is replete with examples that indicate the disputed term is intended to have its plain and ordinary meaning. For example, the specification refers to the second physical host as physical host 160B in Figure 5. Ex. E at 4:4-8, 11:48-55. As is clear from Figure 5, physical host 160A, 160B and 160C are all identical physical host servers, a well-known construct to one of skill in the art. *Id.* at 12:6-15 (“in this example, virtual server 162B is transferred from **old physical host 160A** to **new physical host 160B.**”) (emphasis added).

The remainder of the disputed term (“can accommodate the requested modified resource allocation”) is also used in its plain and ordinary meaning. The claim requires that a request for a modified resource allocation be denied by the first physical host, followed by a determination that a second physical host can accommodate the denied request. This means exactly what one of skill in the art would think it means, that the second physical host must have enough resources to service the previously denied request. This understanding is reflected in the specification,

which notes one exemplary way to evaluate the available host resources is to perform periodic monitoring of the resource loads on each server. *Id.* at 11:12-16.

VMware's proposal that this claim term is indefinite is belied by the claim language itself and the specification as a whole. As the above citations to the intrinsic record show, there is ample evidence that the disclosure relating to the disputed term informs those skilled in the art about the scope of the invention with reasonable certainty. For instance, the resource loads of all of the physical hosts can be monitored in order to have an accurate idea of the resource requests denied by the first physical host, and the available capacity of the remaining physical hosts. *Id.* at 11:12-20. Further, the patentee describes how various embodiments would then use that information to identify a second physical host with sufficient resources. *Id.* at 11:20-35. Further still, the specification describes in detail one exemplary embodiment of the noted easiest fit heuristic, complete with exemplary algorithms. *Id.* at 11:36-57. Therefore, one of skill in the art would be well informed as to what is being claimed and VMware's arguments to the contrary fail to show otherwise and certainly fail to raise to the level of clear and convincing evidence.

VMware's alternative construction should be rejected since it basically a mirror recitation of the disputed term itself. It is also improper because VMware reads out the limitation of "the requested modified resource allocation" and reads in "the denied request to modify a resource allocation." This is unnecessary, however, because the claim language itself makes clear that the disputed "requested modified resource allocation" ultimately refers back to the denied request to modify a resource allocation:

receiving an indication that a first physical host is overloaded, wherein the indication is based on a determination that a virtual server is overloaded and wherein the determination that a virtual server is overloaded is based on one or more resource unavailable messages resulting from denied requests to modify a resource allocation;

determining that a second physical host **can accommodate the requested modified resource allocation** . . .

*Id.* at 14:1-9 (emphasis added). In light of this, if the disputed term needs to be construed as anything other than its plain and ordinary meaning, as VMware contends, it should be construed to incorporate the correct construction<sup>6</sup> of “denied requests to modify a resource allocation,” i.e., “determining that a second physical host can accommodate the request(s) by the virtual server that could not be immediately serviced.” This is exactly the alternate construction that IV has proposed and which the Court should adopt if it is inclined not to apply the plain and ordinary meaning.

8. “determination that a virtual server is overloaded” (’686 patent Claims 5, 6 and 7)

IV’s Proposed Construction	VMWare’s Proposed Construction
plain and ordinary meaning	“determination that an average number of resource denials for a virtual server is beyond a pre-configured threshold”

The disputed term should be given its plain and ordinary meaning and not be further limited to include “an average number of resource denials . . . beyond a pre-configured threshold” as VMware suggests. Review of the intrinsic record reveals that the disputed term is used in its commonly understood manner to those of skill in the art. This is well-illustrated by the disputed term’s constituent parts themselves. First, there can be no real dispute that “determining” is used as that term is commonly understood, i.e., to ascertain, often based on research. Ex. E at 2:62-65 (“monitors resource denials received by virtual servers and determines if a virtual server is overloaded based upon the resource denials.”). Further, as discussed in depth *supra* III.B.5, “virtual server” is used as commonly understood by those of skill in the art. *See, e.g.*, Ex. E at 1:59-2:15; Ex. P at File History, Oct. 5, 2004, Office Action

<sup>6</sup> Refer *supra* section III.B.4 for why this is the correct construction of “denied requests to modify a resource allocation”.



Response, p. 14-15. Finally, “overloaded” in the context of virtual computing is a well-understood construct and is being used as such in the ’686 patent’s claims and specification. *See* Ex. E at claim 5 (“the determination that a virtual server is overloaded is based on one or more resource unavailability messages”) (emphasis added).

VMware’s proposal fails to acknowledge the understanding of one of ordinary skill in the art regarding these basic building blocks of the art itself. More importantly, however, VMware’s proposal significantly limits the disputed term by importing limitations from a lone exemplary embodiment. Presumably, for support, VMware is relying on the portions of the specification which note that “in one embodiment . . . [i]f the resource denials received by a particular virtual server exceed a pre-specified limit, the virtual server is considered overloaded . . .” Ex. E at 3:1-3. Significantly, VMware fails to take into account the broader general description of the disputed term immediately preceding its selected language and within the very same exemplary embodiment. This disclosure shows all that is required with respect to the “determination that a virtual server is overloaded” is that it be based on resource denials. *Id.* at 2:62-65 (“In one embodiment, a dynamic resource configuration module monitors resource denials received by the virtual servers **and determines if a virtual server is overloaded based upon the resource denials.**”) (emphasis added). It is well settled law that “[I]t is improper to read limitations from a preferred embodiment described in the specification—even if it is the only embodiment—into the claims absent a clear indication in the intrinsic record that the patentee intended the claims to be so limited.” *Liebel-Flarsheim Co. v. Medrad, Inc.*, 358 F.3d 898, 913 (Fed. Cir. 2004).

In this instance, as discussed above, the ’686 patent’s specification and claim language do not indicate that VMware’s proposed limitations were intended by the patentee. This conclusion is additionally supported by the prosecution history of the ’686 patent, which indicates that IV’s

interpretation is indeed the full scope of disputed term as intended by the patentee. *See* Ex. E at File History, May 14, 2013 Response to Office Action, p. 5 (“[c]laims 5-7 now recite ‘the determination that a virtual server is overloaded is based on one or more resource unavailability messages resulting from denied requests to modify a resource allocation,’ which Applicants believe would be understood by a person having ordinary skill in the art.”).<sup>7</sup> As discussed above *supra* III.B.3-4, a denied request to modify a resource allocation is a “resource denial.”

Therefore, it is apparent from the intrinsic record as a whole that VMware’s proposal reads in limitations from an exemplary embodiment that neither the patentee nor the patent examiner intended, thus impermissibly limits the claim. Accordingly, VMware’s proposed construction must be rejected.

9. “virtual server overload signal” (’726 patent Claims 1, 4, 5 and 8)

IV’s Proposed Construction	VMWare’s Proposed Construction
“an indication that a virtual server has been or is being denied resources”	“signal indicating that an average number of resource denials for a virtual server is beyond a pre-configured threshold”

As with its proposed construction of the preceding term—and for the same reasons—VMware’s proposal here is impermissibly limiting. IV’s proposal should be adopted because it is consistent with the claim language as well as specification disclosure. As discussed in depth *supra* III.B, the claim language surrounding the disputed term makes it clear that a virtual server overload signal is the result of a resource denial. Therefore, IV’s proposal is most consistent with the claim language and specification disclosure, whereas VMware’s proposal seeks to import limitations into the term in a way never intended by the patentee and contrary to established claim construction law.

<sup>7</sup> The examiner apparently agreed with the patentee’s argument since he immediately transmitted an issue notification based on the patentee’s arguments/amendments.

10. Alleged Means Plus Function Elements ('686 patent Claim 7) ('726 patent Claims 1, 3, 5 and 7)

i. '686 patent, claim 7 “component” terms (i.e., clauses 1-3 of Ex. A)

VMware’s allegations that certain claims of the '686 patent and '726 patent are subject to 35 U.S.C. § 112 ¶ 6 must fail because: (1) “means for” is not present in any of the claims, (2) the prefatory language used does not consist of merely a nonce word followed by function, and (3) the claim language itself provides sufficient structure to avoid the application of § 112 ¶ 6.

First, VMware argues that three of the five elements, of claim 7 of the '686 patent are subject to § 112 ¶ 6. *See* Ex. A at claims 1-3. This is incorrect. The specific elements and claim overall, read in light of the patent’s specification and lack of “means for” language, illustrate that the disputed elements connote sufficient structure to avoid application of § 112 ¶ 6. *See Fisher-Rosemount Sys. v. ABB Ltd.*, 2019 WL 6830806, at \*15-16 (S.D. Tex. Dec. 12, 2019).

Importantly, because the claim does not recite “means for” there is a presumption against applying § 112 ¶ 6. *See Williamson*, 792 F.3d at 1347-49. In order to overcome this presumption, the defendant must demonstrate that the claim terms fail to recite sufficiently definite structure or else recites function without reciting sufficient structure for performing that function. *See Fisher-Rosemount Sys.*, 2019 WL 6830806, at \*16. Neither of these are the case with disputed elements 1-3.

VMware fails to appreciate—or chooses to ignore—the preamble and claim element immediately preceding those elements which it contends are means-plus-function. This error is fatal to VMware’s argument. For example, the preamble of claim 7 of the '686 patent states “a system for modifying the computer resources allocated **to a virtual server operating in a first physical host** of multiple physical hosts, the system comprising.” The next element recites “one

or more **processors** and one or more **memories.**” These are key for interpreting the claim as a whole and help provide structure to the later elements.

Take the recitation in the preamble of “virtual server,” and “operating in a first physical host,” for instance. These two terms are structural in nature. A virtual server is an abstraction of a physical server (i.e., a physical host), and thus both terms collectively provide at least the same structure to one of skill in the art as the term “server” alone. It is well-settled that the term “server” connotes a definite structure to one of skill in the art. *See Sound View Innovations, LLC v. Facebook, Inc.*, 2017 WL 2221177, at \*10-11 (D. Del. May 19, 2017) (“[s]erver has a well-known meaning to a person of ordinary skill in the art and connotes a definite structure.”). The terms “processor” and “memory” are also well-known structural terms. *See, e.g., Fisher-Rosemount Sys.*, 2019 WL 6830806, at \*16 (“processor can on its own recite at least some structure to persons of ordinary skill in the art”) (internal quotations omitted); *Advanced Mktg. Sys., LLC v. CVS Pharmacy, Inc.*, 2016 WL 1741396, at \*20 (E.D. Tex. May 3, 2016) (finding processor and memory together provided sufficient structure to avoid § 112 ¶ 6). These clearly structural terms help clarify that the disputed terms’ use of “component”—generally understood to mean “a constituent part,”<sup>8</sup>—means the remainder of the disputed terms following “component” are parts of the earlier claimed structural elements discussed above, such as, “virtual server,” “physical host,” “processors,” and “memories.” Therefore, the claim language in light of the specification evidences that one of ordinary skill in the art would understand the claimed elements as having a specific structure sufficient to prevent the application of § 112 ¶ 6.

If the Court is inclined to find that the above discussed elements do implicate § 112 ¶ 6, IV proposes the function and structure disclosed in Exhibit A. IV’s proposed structure for

---

<sup>8</sup> *See Netfuel, Inc. v. F5 Networks, Inc.*, 2017 WL 2834538, at \*5-8 (N.D. Ill. June 29, 2017).

elements 1-3 should be adopted based on the specification's disclosures at Figure 1, 2:63-3:4, 3:59-4:8, 5:7-28 and 5:42-62, which disclose that dynamic resource configuration module 100 in combination with physical hosts 160A-C and virtual servers 162A-G determine whether a physical host is overloaded based on resource denials and the process of transferring to a second physical host.

- ii. '726 patent Claims 1, 4, 5, "virtual server resource monitor" and "program code for creating a virtual server resource monitor" terms (i.e., clause 4 of Ex. A)

VMware next asserts that the two elements in clause 4 as identified in Exhibit A are subject to § 112 ¶ 6, however, VMware again is mistaken. With respect to the "virtual server resource monitor" term, the structural disclosures in the preamble help illustrate that one of ordinary skill in the art would understand the structure associated with the disputed element. For example, "network system," "physical hosts," and "virtual server" are all structural terms known in the art as discussed *supra* III.B.10(i), and give context to the environment in which the "virtual server resource monitor" is claimed.<sup>9</sup> In addition, the disputed term itself claims that the "virtual server resource monitor" is "communicatively coupled to the first physical host," two terms that provide well-known structure among those of skill in the art. *See supra* III.B.5.<sup>10</sup>

Furthermore, the specification clearly illustrates that the term "virtual server resource monitor" is not a nonce term, and in fact, has structure that would be apparent to one of skill in the art in light of the specification's disclosure. For instance, the disputed element is a component of "dynamic resource configuration module 100." Ex. F at 4:64-67. "Dynamic resource configuration module 100" is coupled to a physical host machine or resident on a

<sup>9</sup> The same is true for "computer resources." *See also CSB-Sys. Int'l, Inc. v. SAP America, Inc.*, 2011 WL 3240838, at \*12-14 (E.D. Pa. July 28, 2011) (noting when construing server that various dictionaries refer to it as a computer which makes network resources, such as disk drives and printers available).

<sup>10</sup> *See also Collaborative Agreements, LLC v. Adobe Sys. Inc.*, 2015 WL 2250391, at \*12-14 (W.D. Tex. May 12, 2015).

physical host and may be implemented as a software driver. *Id.* at 3:49-56. The disputed element can also be portions of the software code implementing the dynamic resource configuration module, i.e., part of a software driver. *Id.* at 4:67-5:3. Therefore, the “virtual server resource monitor” could be a portion of a software driver that is on, or coupled to, a physical server. *See id.* at Fig. 3. This is more than sufficient for one of skill in the art to know the structure associated with the disputed term and avoids § 112 ¶ 6. *See Intellectual Ventures II LLC v. Bitco Gen. Ins. Corp.*, 2016 WL 125594, at \*8 (E.D. Tex. Jan. 11, 2016) (noting software can connote structure to one of skill in the art and is typically done through outline of an algorithm, flowchart or specific set of rules). Therefore, the disputed “virtual server resource monitor” term does not implicate § 112 ¶ 6 as it has sufficient structure in the claim itself and in the specification.

The second disputed term in this group is similarly situated and does not trigger application of § 112 ¶ 6 for the same reasons discussed above. Furthermore, because the second disputed term in this group includes “program code for creating,” and it has been held that “program code” and similar terms have sufficient structure so as to avoid a means-plus-function classification, this element would also be understood by one of skill in the art as connoting sufficient structure. *See RLIS, Inc. v. Allscripts Healthcare Solutions, Inc.*, 2013 WL 3772472, at \*14 (S.D. Tex. July 16, 2013); *Eolas Techs., Inc. v. Adobe Sys., Inc.*, 810 F. Supp. 2d 795, 810 (E.D. Tex. 2011); *Aloft Media, LLC v. Adobe Sys., Inc.*, 570 F. Supp. 2d 887, 898 (E.D. Tex. 2008).

If the Court is inclined to find that the above discussed elements do implicate § 112 ¶ 6, IV proposes the function and structure disclosed in Exhibit A. IV’s proposed structure for element 4 should be adopted based on the specification’s disclosure at Figure 1, Figure 3, 4:64-

5:4, 5:44-65 which discloses that dynamic resource configuration module 100 includes, among other things, virtual service resource modifier, which in combination with physical hosts 160A-C monitors resource denials and sends a virtual server overload signal. VMware’s attempts to limit the structure to virtual server resource monitor 110, however, fails to account for the claimed functionality as a whole.

- iii. ’726 patent Claims 1, 4, 5, “virtual server resource modifier” and “program code for creating a virtual server resource modifier” terms (i.e., clause 5 of Ex. A)

This group of disputed terms is nearly identical to the two terms discussed immediately above except that instead of claiming a “monitor” they claim a “modifier.” For the same reasons discussed above, these two terms do not implicate § 112 ¶ 6.

First, the structural disclosures in the preamble, such as “network system,” “physical hosts,” “virtual server,” and “computer resources” are all structural terms known in the art and give the necessary context for one of ordinary skill in the art to understand the structure of the terms. *See supra* III.B.10(ii). In addition, the disputed term itself claims that the “virtual server resource modifier” is “communicatively coupled to the first physical host,” two terms that provide specific structure to those of ordinary skill in the art. *Id.*

Similar to the “monitor” terms, in this case, the specification strongly indicates that one of skill in the art would know the structure of the disputed terms. The disputed element is a component of “dynamic resource configuration module 100,” that in turn is coupled to a physical host machine or resident on a physical host and may be implemented as a software driver, of which the disputed element can be a part. Thus, just like the “virtual server resource monitor” terms, the disputed terms here could be a portion of a software driver that is on, or coupled to, a

physical server, which taken collectively are sufficient for one of skill in the art to know the structure associated with the disputed term and avoid § 112 ¶ 6. *See supra* III.B.10(ii).<sup>11</sup>

The second disputed term in this group “program code for creating a virtual server resource modifier” also does not trigger application of § 112 ¶ 6 for the same reasons discussed above with respect to “virtual server resource modifier.” *Id.* As with the other “program code” elements, additional evidence that the application of means-plus-function to this element is incorrect is found in the fact that courts routinely hold “program code” and similar terms have been held to provide enough structure to avoid § 112 ¶ 6. *See RLIS, Inc.*, 2013 WL 3772472, at \*14; *Eolas Techs., Inc.*, 810 F. Supp. 2d at 810; *Aloft Media, LLC*, 570 F. Supp. 2d at 898. As a result, when combined with the evidence set forth above, the disputed “program code” term would clearly be understood by one of skill in the art as connoting sufficient structure to avoid means-plus-function treatment.

If the Court is inclined to find that the above discussed elements do implicate § 112 ¶ 6, IV proposes the function and structure disclosed in Exhibit A. IV’s proposed structure for element 5 should be adopted based on the specification’s disclosure at Figure 1, Figure 2A, Figure 4, 4:64-5:20, which discloses that dynamic resource configuration module 100 includes, among other things, virtual service resource modifier, and in combination with physical hosts 160A-C and virtual servers 162A-G receives virtual server overload signals and signals a resource modification is needed. VMware’s attempts to limit the structure to virtual server resource monitor 120, however, fails to account for the claimed functionality as a whole as more than just the monitor 120 is necessary.

---

<sup>11</sup> *See also* Ex. F at Fig. 4 (illustrating exemplary flow chart for the disputed term).



- iv. '726 patent Claims 1, 4, 5, “a load balance[ing / er] module” and “program code for creating a load balancing module” terms (i.e., clause 6 of Ex. A)

VMware’s argument that the disputed terms “load balancing module” and “program code” for creating the same are means-plus-function elements is particularly illogical given that the term is a specific and well-known construct with a structure known by those of skill in the art. This is reflected in the plain language of the claim, the specification of the '726 patent, and as apparent from the file history.

The term itself (load balancing module) is structural. A load balancer, in the context of virtual computing is exactly what it sounds like, a module that balances the resource load across various physical hosts. As is described in the patent’s specification, this is typically implemented in software on or communicatively coupled, to one or more of the physical hosts. Ex. F at 4:67-5:3; 3:49-56. Not surprisingly, this is also how the claimed load balancing module is described in the claims themselves—communicatively coupled to the plurality of physical hosts. This is further illustrated in Figure 5, which provides an exemplary flow chart and schematic illustrating the structure of the load balancing software module, and in column 11 lines 22-52 where one embodiment of a load balancing algorithm is shown.<sup>12</sup> Therefore, the disputed term does not implicate § 112 ¶ 6, as it has sufficient structure in the claim itself in light of the specification.<sup>13</sup>

If the Court is inclined to find that the above discussed elements do implicate § 112 ¶ 6, IV proposes the function and structure disclosed in Exhibit A. IV’s proposed structure for

<sup>12</sup> See *Intellectual Ventures II LLC*, 2016 WL 125594, at \*8 (software structure typically shown via an outline of an algorithm, flowchart or specific set of rules).

<sup>13</sup> The same holds true for the second element in this group—program code for creating a load balancing module—with the addition of the established claim construction principal discussed above stating that the term “program code” and similar terms do not implicate § 112 ¶ 6. See *RLIS, Inc.*, 2013 WL 3772472, at \*14; *Eolas Techs., Inc.*, 810 F. Supp. 2d at 810; *Aloft Media, LLC*, 570 F. Supp. 2d at 898.

element 6 should be adopted based on the specification's disclosure at Figure 1, Figure 5, Figure 6, 4:64-5:20, 10:53-11:20, which discloses that dynamic resource configuration module 100 includes, among other things, load balancer, and in combination with physical hosts 160A-C and virtual servers 162A-G receives virtual server resource modification signal and determines whether the physical host is overloaded and sends a host transfer signal if it is. VMware's attempts to limit the structure to physical host load balancing module 130, however, fails to account for the claimed functionality as a whole, since as seen in the above citations, more than just the load balancer 130 is required.

- v. '726 patent Claims 1, 4, 5, "dynamic virtual server mover" and "program code for creating a dynamic virtual server mover" terms (i.e., clause 7 of Ex. A)

This group of disputed terms is nearly identical to the virtual server monitor and virtual server modifier terms discussed *supra* III.B.10(ii)-10(iii), however, relate to a virtual server mover. For the same reasons discussed above in Sections 10(ii)-10(iii), these two terms do not implicate § 112 ¶ 6. Specifically, the same structural language appears in the claim and the same portions of the specification address the virtual server mover. It would be clear to one of skill in the art that the virtual server mover is described in structural terms as software coupled to or running on one or more physical host servers, which is itself structural in nature and well-understood. *See supra* III.B.10(ii)-(iii). Further, the "dynamic virtual server mover" is "communicatively coupled to the plurality of physical hosts," providing further structure. *Id.*<sup>14</sup> Therefore, for the same reasons as discussed *supra* Section 10(ii)-10(iii), application of § 112 ¶ 6 is not required.

---

<sup>14</sup> *See also* Ex. F at Fig. 6 (exemplary flow chart).

If the Court is inclined to find that the above discussed elements do implicate § 112 ¶ 6, IV proposes the function and structure disclosed in Exhibit A. IV's proposed structure for element 7 should be adopted based on the specification's disclosures at Figure 1, Figure 6, 4:64-5:20, 6:20-31, 6:39-59, 12:1-10, which disclose that dynamic resource configuration module 100 includes, among other things, dynamic virtual service mover and in combination with physical hosts 160A-C, receives the physical host transfer signal. VMware's attempts to limit the structure to dynamic virtual server mover 140, however, fails to account for the claimed functionality as a whole as more than just the mover 140 is necessary.

- vi. '726 patent Claims 3, 7, "dynamic virtual server mover further configured to" term (i.e., clause 8 of Ex. A)

The same arguments made above with respect to claims 1, 4 and 5 (*supra* Sections 10(ii), 10(iii) and 10(v)) hold true for this claim as well. The same specification support is applicable and same law indicates that this is not a means-plus-function claim. IV will not re-hash these arguments here for the sake of brevity and judicial efficiency, however, if the Court is inclined to find that this term is subject to § 112 ¶ 6, IV proposes the function and structure disclosed in Exhibit A. IV's proposed structure for element 8 should be adopted based on the same evidence and reasoning cited *supra* Section 10(v).

### **C. The '818 Patent**

Claiming priority to December 7, 2007, the '818 patent discloses a set of inventions that enable more granular quality of service ("QoS") to be provided in the context of a virtual input/output ("I/O") environment.

- 1. "virtual [network/storage network] interface layer of an application server" (Claims 1, 17, 32, 42)

IV's Proposed Construction	VMware's Proposed Construction
Plain and ordinary meaning	“a virtual storage network interface to higher layers of the virtual node in an application server” / “virtual network interface layer to higher layers of the virtual node in an application server” / “virtual interface layer to higher layers of the virtual node in an application server”

The claim limitations “virtual [network/storage network] interface layer of an application server” should be given their plain and ordinary meaning. The term’s use in the specification and throughout the claims is consistent with its customary use as would be evident to one of skill in the art at the time of the invention. For example, the specification describes examples of virtual interface layers (e.g., virtual network interface 220, virtual HBA 208a) that “emulate” layers of a networking or storage protocol stack using, for example, encapsulation techniques. *See* Ex. G at 3:67-4:62, 7:40-43, Fig. 2 (describing “protocol stack components and modules,” “virtual interface drivers,” and associated interface layers in exemplary embodiment). The intrinsic record thus describes virtual interface layers in a manner that is fully consistent with how they are understood in the art.

VMware’s proposed inclusion of the phrase “to higher layers of the virtual node” would unnecessarily qualify the “virtual interface layer of an application server” recited in the claim. The proposed language would read out embodiments disclosed in the specification by suggesting that a virtual interface layer can interface with multiple higher layers, when in fact the specification shows that a virtual layer, in at least some instances, can interface with a single higher layer. *See, e.g., id.* at 4:14-32; 4:51-62; Fig. 2 (e.g., layer 208a interfacing with higher layer 209, layer 220 interfacing with higher layer 222, etc.).

The proposed inclusion of the additional “virtual node” qualifier is similarly extraneous and contrary to the literal expression of the claim language. Notably, the term “virtual node” (as “virtual node identifier”) appears elsewhere in claims 32 and 42, but does not appear at in claims 1 or 17, indicating that the patentee knew how to use the term in claim drafting and intentionally chose not to recite it in certain claim elements. The specification includes no limitations that would otherwise impose a “virtual node” requirement on the proposed terms. *See Cont’l Circuits LLC v. Intel Corp.*, 915 F.3d 788, 796-97 (Fed. Cir. 2019) (specification must contain expressions of manifest exclusion or restriction representing a clear disavowal to limit claim scope).

2. “hierarchical token bucket resource allocation” / “token(s)” (Claims; 1, 17, 30, 32, 33, 37-42)

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	The specific class-based scheduling algorithm known in the art as the “hierarchical token bucket” / “token” as used in “hierarchical token bucket resource allocation”

The term “hierarchical token bucket resource allocation” should be given its plain and ordinary meaning, as understood by persons having ordinary skill in the art and supported by the specification. The specification explains that the claimed invention involves “a two-tier hierarchical QoS management process . . . employed in a virtual I/O server” in which “network fabric resources are allocated in a hierarchical arrangement.” *See* Ex. G at 2:3-18. Said hierarchical resource allocation can be implemented “using scheduling and queuing methods such as hierarchal token bucket.” *Id.* at 9:61-65. The specification then goes on at length to describe exactly how the claimed resource allocation scheme is achieved using hierarchical token buckets in a non-limiting manner. For example, the specification explains that “HTB includes hierarchical classes where three class types exist: root, non-leaf and leaf” and further:

HTB mechanisms allocate so-called tokens for the buckets at regular intervals. Scheduling a message or packet for transmission results in deducting an amount of tokens from a corresponding bucket, and is permitted when the corresponding bucket includes a sufficient number of tokens

*See id.* at 9:61-10:67, Fig. 2, Figs. 12-13. Accordingly, there is no need to construe the term “hierarchical token bucket resource allocation” beyond its plain and ordinary meaning since the intrinsic record unambiguously conveys the intended meaning of the term and does so in a manner that is entirely consistent with the well-understood meanings of the constituent terms “hierarchical token bucket” and “resource allocation.”

VMware’s proposed construction, to the extent it is ascertainable, cites to a series of extrinsic papers that teach various implementations of a hierarchical token bucket algorithm as used in the art. *See* Ex. Q (disclosing Linux implementations, a WLAN implementation, 802.11 implementation, etc.). This runs contrary to the specification’s non-limiting disclosure. Indeed, VMware’s proposal, that multiple different implementations of a hierarchical token bucket algorithm were known, supports the fact the potential applications of hierarchical token bucket are numerous and should not be limited to a specific, extrinsic definition as VMware proposes. VMware’s construction further ignores the significance of “. . . resource allocation [of bandwidth],” which qualifies the “hierarchical token bucket” language and applies it to the specific context of the claimed invention. *See, e.g., Continental Circuits, LLC*, 915 F.3d at 799 (while extrinsic evidence can shed useful light on the relevant art, it is less significant than the intrinsic record in determining legally operative meaning of disputed claim language). *See also* Ex. G, 10:15-33 (describing implementation in context of allocating bandwidth).

The standalone term “token” is similarly unambiguous and extensively discussed in the specification in a manner that is fully consistent with “hierarchical token bucket resource

allocation” and its usage in the art. For instance, “tokens” are disclosed as being used to “schedule and shape traffic . . . [where] each class or node in the hierarchy has a bucket of tokens associated with it . . . [and scheduling a packet for transmission] results in the deducting an amount of tokens from a corresponding bucket.” Ex. G at 10:23-28. Thus, because the specification discloses the use of “token” in its common parlance it should also be given its plain and ordinary meaning.

3. “enforcing . . .” / “receiv[e/ing] . . .” / “classify[ing] . . .” / “compar[e/ing] . . .” / “forward[ing] . . .” / “buffer[ing] . . .” (Claims; 1, 17, 30, 32, 37-39, 42)

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	“enforcing . . . across the physical [storage network] interface of the virtual I/O server” / “receiv[e/ing] in the virtual I/O server” / “classify[ing] in the virtual I/O server” / “compar[e/ing] in the virtual I/O server” / “forward[ing] in the virtual I/O server” / “buffer[ing] in the virtual I/O server”

As an initial matter, VMware proposes to construe the leading verb of almost all the disputed terms to read in the phrase “in the virtual I/O server.” This is impermissible, however, because each of the disputed terms are well-understood computer networking terms that are used in the specification in a manner entirely consistent with their ordinary and customary usage. *See, e.g.*, Ex. G at 11:38-12:46, Fig. 5 (exemplary embodiment that identifies and discusses each of the terms at issue). The terms should be given their plain and ordinary meaning.

VMware’s attempt to insert the phrase “in the virtual I/O server” to qualify nearly every element recited in the claims at issue when the term is only found in the preamble is improper. *See Catalina Mktg. Int’l, Inc. v. Coolsavings.com, Inc.*, 289 F.3d 801, 808 (Fed. Cir. 2002) (facially structural preamble limitation did not limit claim body where not recited in body, only limited separate claim in which it appeared in claim body as well). Here, the claim drafters

defined a structurally complete invention and were well aware of the term “virtual I/O server” as evidenced by their decision to expressly include said term in the preamble of some independent claims (e.g., 32, 42) and not others (e.g., claim 1), and not to include said term in the body of any of the disputed claims.

VMware’s proposed constructions also contradict the specification and would, thus, read preferred embodiments out of the claim. The specification explains that “virtual I/O servers can be used to create logical separations between the application servers and I/O subsystems to make the I/O subsystems as logical resource units to application servers” and “allow for use of existing computing infrastructures, including hardware and software, while abstracting the operation of the intermediate I/O switch fabric.” *See, e.g.*, Ex. G at 1:31-36; 4:14-23 (explaining how “virtual I/O server” is logically, but not necessarily physically distinct from the application server(s) and I/O subsystems that it connects). By reading in a “virtual I/O server” limitation into almost every claim element, VMware is improperly attempting to redefine structural relationships between their proposed virtual I/O server limitation and other structural claim elements such as “physical [network/storage network] interface(s),” “local area network packets,” and “virtual [network/storage network] interface layer.” These proposed limitations run contrary to both the plain meaning of the claim language and the specification’s characterization of a virtual I/O server as a virtual component that is logically, but not necessarily physically, abstracted from underlying server and I/O subsystems.

4. “maintain[ing] a connection over a network fabric” (Claims 1, 17, 30, 32, 42)

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	"maintaining a connection between the physical interface of the application server and the physical interface of the virtual I/O server over a network fabric"



The term “maintain[ing] a connection over a network fabric” should be given its plain and ordinary meaning. The specification makes it clear that the disputed term is being used in a manner commonly understood by those of skill in the art. For example, the specification explains that “[i]n one implementation, the virtual network interface 220 is used to access network interfaces of the virtual I/O server 106 over the I/O switch fabric, using the encapsulation module 206 to provide the interfaces to establish and maintain the connection.” Ex. G at 5:17-21; Fig. 1 (illustrating fabric 104; connections to application servers 102 and virtual I/O server 106). The intrinsic record provides no basis to deviate from the well-understood meaning of this term.

VMware’s proposed construction seeks to construe the first element of every independent claim at issue (the “maintain[ing] . . .” element) to introduce additional structural limitations that are neither expressed in the claim language nor suggested by the specification. Like the proposed constructions discussed above in Section III.C.3, VMware is attempting to insert additional structural limitations “physical interface of the application server” and “physical interface of the virtual I/O server” that are not present in the relevant claims. Such language is unnecessary and would improperly import limitations from the specification and preamble into the body of the claims. *See Continental Circuits LLC*, 915 F.3d at 796-97 (cautioning that specification must clearly and manifestly disavow claim scope to have limiting effect). *See also Retractable Techs., Inc. v. Becton, Dickinson & Co.*, 653 F.3d 1296, 1306 (Fed. Cir. 2011); *Epistar Corp. v. Int’l Trade Comm’n*, 566 F.3d 1321, 1335 (Fed. Cir. 2009). In addition, the claim drafters were well aware of the terms “virtual I/O server” (recited in preamble), “application server” (recited elsewhere in claims), and “physical [interface]” (recited elsewhere

in claims), and consciously chose not to include these known terms in the claim elements at issue.

VMware’s proposed language also runs contrary to the specification, which teaches that a virtual I/O server need not be physically distinct from the application servers and associated virtual network interface(s). *See, e.g.*, Ex. G at 4:9-13 (“virtual network interface, in one implementation, emulates an Ethernet NIC [and] plugs in at the bottom of the network stack and provides an Internet Protocol address bridged by the Virtual I/O server”). In this regard, VMware’s proposed constructions improperly seek to read additional physical, structural limitations into the claims (e.g., additional “physical interface(s)” relative to the “application server(s)” and the proposed “virtual I/O server”) that are neither literally expressed in the claim language nor required by the specification.

5. Alleged Means Plus Function Elements (Claim 17) (individually set forth in Ex. C)<sup>15</sup>

As an initial matter, VMware seeks to construe a large swath of claim 17 as subject to § 112 ¶ 6 and indefinite. However, claim 17 is not subject to § 112 ¶ 6 for at least the following reasons. First, the lack of the term “means” in claim 17 raises a presumption that the claim limitations at issue are not means-plus-function limitations. *Watts* at 881. Second, to overcome this presumption, VMware is required to demonstrate that “the claim term fails to recite sufficiently definite structure or else recites function without reciting sufficient structure for performing that function.” *Inventio AG v. ThyssenKrupp Elevator Am. Corp.*, 649 F.3d 1350, 1356 (Fed. Cir. 2011) (internal quotations/citations omitted). But the claims themselves tie any recited function to sufficiently definite structure. For example, with respect to element 1 of Ex. C, the limitation “maintain a connection, over a network fabric, to a virtual storage network

---

<sup>15</sup> For ease of reference we have reproduced the parties proposed constructions of these 10 terms in Exhibit C.

interface layer of an application server, wherein the virtual storage network interface layer is associated with a virtual storage node identifier” distinctly ties the function of “maintain[ing] a connection . . . to a virtual storage network interface layer of an application server” to its structure, including a network fabric, a virtual storage network interface layer, and a virtual storage node identifier. Elements 2-3 similarly recite the limitation “physical storage network interface,” which is well-understood by those skilled in the art to be structural. *See* Ex. G at 3:4-7; Fig. 1 (denoting exemplary structure associated with “physical storage network interface”). With respect to remaining elements 4-10 of Ex. C, each disputed element either individually recites structure or relates back to the structural limitations discussed above. Ex. G at claim 17 (“connection,” “virtual storage network interface layer,” “the storage command,” “the current amount of tokens,” “destination” (elements 4-10 of Ex. C)). *See, e.g.,* Ex. G at 4:6-9; 5:10-15; 7:28-31; 11:18-22; 12:7-12; 12:41-46; 13:6-17; 16:6-19; Figs. 1-2; Figs. 4-10.

Further, claim 17 includes additional elements external to the disputed elements that convey structure such as “input/output fabric interface,” “storage network interface,” and “a memory.” These additional elements, read in conjunction with the disputed claim language and considered in their entirety indicate the presence of structure to a PHOSITA.

It is only once the established presumption is rebutted—i.e., that it has been shown that the disputed term(s) do not connote sufficiently definite structure within the claims, *and* that there’s a lack of corresponding structure within the specification—that the disputed terms could be deemed indefinite. *See Apple Inc. v. Motorola, Inc.*, 757 F.3d 1286, 1296-97 (Fed. Cir. 2014). Here, however, VMware cannot meet this burden because, as discussed above, the disputed terms include sufficient structure within the claims themselves read in light of the specification.

Further, for VMare’s argument to succeed, VMware would be required to show that “one or more input/output virtualization modules” is a nonce term that would not be understood by a PHOSITA as denoting structure, which it is not. *See Zeroclick, LLC v. Apple Inc.*, 891 F.3d 1003, 1007-09 (Fed. Cir. 2018). The term module is qualified by “input/output,” which denotes structure as understood by a PHOSITA. *Id.* (refusing to apply § 112 ¶ 6 where alleged nonce term “code” was qualified as “user interface code”). *C.f. Williamson* at 1351 (applying § 112 ¶ 6 where language “distributed learning control” was insufficient to connote structure onto phrase “distributed learning control module”).

If the Court is inclined to find that the above discussed elements do implicate § 112 ¶ 6, IV proposes the function and structure disclosed in Ex. C. As indicated therein, structure for these elements should be adopted based on the specification’s disclosure of at least 2:9-18; 3:9-11; 3:15-30; 3:43-53; 3:60-4:13; 4:27-5:32, 5:55-6:42, 7:26-67; 8:1-12:28; 12:30-13:4; 13:6-14:29; 14:65-15:23; 15:46-17:19; Figs. 1-4; and Figs. 11-13; describing the structure associated with application server 102, I/O switch fabric 104, virtual I/O server 106, HBA 108, I/O fabric interfaces 110 and 202, fabric driver stack 204, encapsulation module 206, virtual HBA 208a, and virtual network interface 220.<sup>16</sup>

#### **D. The ’051 Patent**

Claiming priority to an application filed on March 15, 2000, the ’051 patent discloses a set of inventions that enable remote service providers to provide private, overlapping address spaces to multiple customers in a virtualized server environment.

1. “customer forwarding [table(s)/information]” (claims: 1, 3)

---

<sup>16</sup> See also Ex. H at Figs. 5-10 (flowcharts illustrating various implementations of the structure at issue).

IV's Proposed Construction	VMware's Proposed Construction
"table(s) containing [a set/sets] of customer specific forwarding information" / "set(s) of customer specific forwarding information"	See construction of "storing a customer lookup table, the customer lookup table storing associations between physical interfaces and tunnel identifiers identifying tunnels for private networks and a plurality of customer forwarding tables" / "storing . . . customer forwarding information, . . . the customer forwarding information associating network addresses with physical interfaces and tunnel identifiers"

IV's proposed constructions of "customer forwarding [table(s)/information]" would add the words "specific" and "set(s)" to further clarify that each individual instance of "customer forwarding table" or "customer forwarding information" should be understood as a distinct piece of information relative to the remaining "plurality" of customer forwarding tables and "multiple" customer forwarding information recited elsewhere in the claim. IV's proposed language is narrowly tailored to track both the underlying claim language and the intended meaning of the terms as expressed in the specification.

Regarding the proposed word "specific," the specification explains that "the information in the customer forwarding tables is segregated by customer because the private address spaces of different customers may overlap . . . ." Ex. H at 12:21-26 (emphasis added). *See also id.* at 11:50-11:58; 11:64-12:3; 12:60-13:7; Fig. 7; Fig. 9 (illustrating separation between sets of customer forwarding tables/information 910, 920, 930, etc.). As a result of this segregation, there is a single, "correct" customer forwarding table/information accessible to each customer from amongst the plurality of customer forwarding tables/set of multiple customer forwarding information. *Id.* at 11:52-54; 12:64-13:11; Fig. 9 (emphasis added). Thus, IV is proposing the phrase "customer specific" to help explain that each individual instance of "customer forwarding

table” and “customer forwarding information” is a distinct piece of information associated with a particular customer.

Regarding the proposed words “set” / “sets,” the above reasoning also applies insofar as each instance of “customer forwarding table” and “customer forwarding information” is defined as a logically distinct or “segregated” data structure. *Id.* The proposed phrasing “[set/sets] of customer specific forwarding information” further clarifies that the “correct” customer forwarding table/information is a logically distinct piece of information relative to the remaining instances of customer forwarding tables/information (the “plurality” / “multiple”).

In addition, the term “information” is an uncountable noun having an irregular plural form such that it is not grammatically correct to state “informations;” rather the proper plural phrasing is e.g., “pieces of information.” The proposed inclusion of the words “set/sets” clarifies the grammatical construct consistently with its usage in the claim.

VMware’s proposed constructions refer to its respective proposed constructions for the entire “storing . . .” element recited in claims 1 and 3. *See infra* III.D.5. In both instances, VMware proposes to modify the language of these disputed terms to exclude the term “customer,” and selectively exclude the term “forwarding,” where they precede and qualify the terms “table(s)” / “information.” It is unclear how removing unambiguous and important claim terms, such as “customer” and “forwarding,” serves to clarify the meaning(s) of the disputed terms in any way. Reading out meaningful claim terms and otherwise failing to provide any rationale for doing so is inconsistent with well-established principles of claim construction. *See, e.g., Bell Commc’ns. Research, Inc. v. Vitalink Commc’ns. Corp.*, 55 F.3d 615, 619-20 (Fed. Cir. 1995) (internal citations omitted) (“the language of the claim defines the scope of the protected invention . . . resort must be had in the first instance to the words of the claim”).

## 2. “virtual server(s)” (claims; 1, 3, 6)

IV’s Proposed Construction	VMware’s Proposed Construction
<p>Plain and ordinary meaning</p> <p>Alternatively – “virtual machine(s) that reside(s) on a physical server and use(s) the physical server’s resources, but [has/have] the appearance of being a separate, dedicated machine(s)”</p>	<p>“a process executing on a host computer that accept [sic] communications requests”</p>

The specification supports construing “virtual server” with its plain and ordinary meaning. For example, the specification explains how:

[A] server application executing on a single physical host can be programmed to process requests made to multiple network addresses. Such functionality is known as virtual hosting. In virtual hosting . . . the virtual host server can service requests to multiple network addresses or domain names. Thus, the functionality of numerous hosts is provided by a single physical host computer, servicing requests made to a plurality of network addresses and domain names by multiple customers.

Ex. H at 2:50-64 (emphasis added). This disclosure demonstrates that the patentee was using the term in a non-limiting manner.

If it is determined that “virtual server(s)” must be construed beyond its plain and ordinary meaning, IV proposes an alternate construction consistent with the specification that accounts for the context of surrounding claim language. For example, in claim 1 the term “virtual servers” is recited in the context of “a host computer containing a plurality of virtual servers which support a private network address space wherein the private network address spaces of two or more of the virtual servers overlap.” The context for claim 3 is similar. The specification explains that a “virtual host server can service requests to multiple network addresses or domain names. Thus, the functionality of numerous hosts is provided by a single physical host computer . . .” Ex. H at 2:60-64 (emphasis added). *See also id.* at Fig. 6 (illustrating logical separation/abstraction of each virtual server 660, 680 in relation to respective physical hosts 640, 650).

In other words, a virtual server is something that appears to be an ordinary server from the outside, but internally it is functionally isolated from its underlying physical host and/or other virtual machines or applications running thereon. IV’s proposed alternate construction also closely tracks well established extrinsic definitions of “virtual server.” *See, e.g.*, Ex. O at 555.

VMware’s proposed construction fails to distinguish a virtual server from a conventional physical server in any meaningful way. For example, VMware’s construction of “a process executing on a host computer that accept[s] communication requests,” is inherently limiting the disputed term to “a process” running on a traditional physical server (i.e. “host computer”). Because the patentee is claiming a virtual server, i.e. a software abstraction of physical hardware operating to appear to a user as a traditional server, the specification’s disclosure is broader than VMware argues. *See, e.g.*, Ex. H at 2:60-64, Fig. 6.

### 3. “physical interface(s)” (claims; 1, 3)

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	“hardware that provides a point of communication between two or more devices”

The disputed term “physical interface(s)” should be given its plain and ordinary meaning. The term’s use in the specification and throughout the claims is consistent with its customary use as would be evident to one of skill in the art at the time of the invention. For instance, the specification discloses “physical interface” consistent with the understanding of one of skill in the art. *See* Ex. H at 10:37-42. The intrinsic record does not otherwise limit the plain and ordinary meaning of this term.

Besides there being no need to construe the term “physical interface(s),” VMware seeks a construction that imports the additional limitation that a physical interface be “between two or more devices.” For example, the specification teaches that “[a] tunnel switch comprises one or more physical interfaces . . .” *See* Ex. H at 10:37-38. Thus, where portions of the specification



clearly make reference to a single component having one or more physical interface(s)– e.g., a tunnel switch –VMware’s proposed construction is improper because it attempts to affirmatively read in a limitation “between two or more devices,” which is not a necessary attribute of a physical interface.

4. “storing . . .” / “determining . . .” / “using . . .” (claims; 1, 3)
  - i. “storing a customer lookup table, the customer lookup table storing associations between physical interfaces and tunnel identifiers identifying tunnels for private networks and a plurality of customer forwarding tables” / “storing customer lookup information . . . the customer lookup information specifying associations between physical interfaces and tunnel identifiers identifying tunnels for private networks and multiple customer forwarding tables”

IV’s Proposed Construction	VMware’s Proposed Construction
<p>Not a claim “term” for purposes of construction;</p> <p>Alternatively – plain and ordinary meaning;</p> <p>Alternatively – “storing a customer lookup table having associations between physical interfaces and first tunnel identifiers that identify tunnels for private networks and identify a plurality of customer forwarding tables” /</p> <p>“storing customer lookup information specifying associations between physical interfaces and first tunnel identifiers identifying tunnels for private networks”</p>	<p>“storing a table that associates a customer identifier with a pair of an incoming physical interface and an incoming tunnel identifier” /</p> <p>“storing information that associates a customer identifier with a pair of an incoming physical interface and an incoming tunnel identifier”</p>

- ii. “storing a plurality of customer forwarding tables, the customer forwarding tables associating network addresses with physical interfaces and tunnel identifiers” / “storing . . . customer forwarding information . . . the customer forwarding information associating network addresses with physical interfaces and tunnel identifiers”

IV's Proposed Construction	VMware's Proposed Construction
<p>Not a claim "term" for purposes of construction;</p> <p>Alternatively – plain and ordinary meaning;</p> <p>Alternatively – “storing tables containing sets of customer specific forwarding information that each associate network addresses with physical interfaces and second tunnel identifiers” /</p> <p>“storing sets of customer specific forwarding information each associating network addresses with physical interfaces and second tunnel identifiers”</p>	<p>“storing a plurality of tables that each associate a network address with a pair of an outgoing physical interface and an outgoing tunnel identifier” /</p> <p>“storing information that associates a network address with a pair of an outgoing physical interface and an outgoing tunnel identifier”</p>

- iii. “determining the correct customer forwarding table from the customer lookup table using the physical interface and the tunnel identifier” / “determining the correct customer forwarding information from the customer lookup information using the physical interface identifier and the tunnel identifier”

IV's Proposed Construction	VMware's Proposed Construction
<p>Not a claim "term" for purposes of construction;</p> <p>Alternatively – plain and ordinary meaning;</p> <p>Alternatively – “determining the correct table with the correct set of customer specific forwarding information from the customer lookup table using the physical interface and the first tunnel identifier” /</p> <p>“determining the correct set of customer specific forwarding information from the customer lookup information using the physical interface identifier and the first tunnel identifier”</p>	<p>“determining the correct customer forwarding table from the customer lookup table using the incoming physical interface and the incoming tunnel identifier” /</p> <p>“determining the correct customer forwarding information from the customer lookup information using the incoming physical interface identifier and the incoming tunnel identifier”</p>

- iv. “determining via the customer forwarding table a physical interface and tunnel identifier associated with a network address of the transmission” / “using the customer forwarding information to identify a physical interface and tunnel identifier associated with a network address of the transmission”

IV's Proposed Construction	VMware's Proposed Construction
<p>Not a claim “term” for purposes of construction;</p> <p>Alternatively – plain and ordinary meaning;</p> <p>Alternatively – “determining, via the correct table of customer specific forwarding information, a physical interface and second tunnel identifier associated with a network address of the transmission” /</p> <p>“using the correct set of customer specific forwarding information to identify a physical interface and second tunnel identifier associated with a network address of the transmission”</p>	<p>“determining via the customer forwarding table an outgoing physical interface and an outgoing tunnel identifier associated with a network address of the transmission” /</p> <p>“using the customer forwarding information to identify an outgoing physical interface and an outgoing tunnel identifier associated with a network address of the transmission”</p>

IV does not view these entire claim elements as claim “terms” requiring a construction for the purposes of claim construction. *See, e.g., O2 Micro Int’l Ltd. v. Beyond Innovation Tech. Co.*, 521 F.3d 1351, 1362 (Fed. Cir. 2008) (district courts not required to construe every limitation present in patent's asserted claims and must only interpret scope of claim terms for which the parties have presented a fundamental dispute). *See also Acumed LLC v. Stryker Corp.*, 483 F.3d 800, 806 (Fed. Cir. 2007) (claim construction does not require purging every shred of ambiguity). However, to the extent that any of these proposed constructions are warranted, the appropriate construction is plain and ordinary meaning.

In addition, the specific constructions of key constituent claim terms found within VMware’s proposals—such as “customer forwarding [table(s)/information]” and “physical interface(s)” —are separately at issue and more appropriately focus on the specific terms. No additional claim language is ambiguous or unclear in light of the specification, which fully describes each and every element of claims 1 and 3. *See, e.g., Ex. H at 10:61-11:41; Figs. 6-7* (providing detailed walkthrough of tunnel switching embodiments as implemented in claimed private virtual networking environment).

In addition to extending beyond the usual scope of claim construction, VMware’s proposed constructions seek to insert the words “incoming,” “outgoing,” and “pair,” to qualify the terms “tunnel identifier(s)” and “physical interface(s)”. The proposed inclusion of the terms “incoming” and “outgoing” would improperly import limitations from the specification into the claim. *See Continental Circuits; Retractable Techs; Epistar Corp. supra* III.C.4. For instance, the specification contemplates the ability of a tunnel switch to switch traffic across sets of tunnels in both directions. Ex. H at 11:51-54; 12:6-9; 12:17-20; 14:20-23 (“tables 800 and 900 operate to switch transmissions in both directions” . . . “the tunnel functions as a bi-directional data pipe”) (emphasis added). Thus, contrary to VMware’s attempt to limit the disputed terms to a single directionality requirement, the patent contemplates both unidirectional and bidirectional embodiments.

VMware’s further proposed inclusion of the word “pair” to qualify the first and second elements of claim 1 and the first element of claim 3 is similarly improper and unsupported by the intrinsic record. The claim language at issue already uses the conjunctive term “and” to group “physical interface(s)” and “tunnel identifier(s)” when intended. Inserting an additional “pair” requirement, as VMware suggests, could improperly mislead a PHOSITA to believe, for example, that the claims require tunnel identifier and physical interface information to be grouped or processed in a certain manner (e.g., contemporaneously, sequentially, etc.) not required by the specification.

The specification also clarifies that “a tunnel switch comprises one or more physical interfaces, with each interface capable of carrying many multiplexed tunnels.” *Id.* at 10:37-39 (emphasis added). VMware’s proposed constructions would contradict these teachings by reading in an artificial distinction between “incoming physical interface” and “outgoing physical

interface,” which could improperly lead a PHOSITA to believe that more than one physical interface is required, contrary to the express language in the specification.

If any constructions beyond plain and ordinary meaning are necessary, IV’s proposed alternate construction would insert the words “first” and “second” to differentiate between separate instances of tunnel identifier(s), but otherwise would leave the intended meaning of the claim language intact without importing VMware’s extraneous limitations.

#### **IV. Conclusion**

For the reasons stated herein IV’s proposed constructions should be adopted.

March 6, 2020

Respectfully submitted,

/s/ Robert Gilman

DEREK GILLILAND  
State Bar No. 24007239  
SOREY, GILLILAND & HULL, LLP  
P.O. Box 4203  
109 W. Tyler St.  
Longview, Texas 75601  
903.212.2822 (telephone)  
derek@soreylaw.com

KARL RUPP  
State Bar No. 24035243  
NIX PATTERSON L.L.P.  
1845 Woodall Rodgers Fwy., Suite 1050  
Dallas, Texas 75201  
972.831.1188 (telephone)  
972.444.0716 (facsimile)  
krupp@nixlaw.com

OF COUNSEL:

Paul J. Hayes  
phayes@princelobel.com  
Matthew Vella  
mvella@princelobel.com  
Robert R. Gilman  
rgilman@princelobel.com  
Jonathan DeBlois  
jdeblois@princelobel.com  
Alex Breger  
abreger@princelobel.com  
PRINCE LOBEL TYE LLP  
One International Place, Suite 3700  
Boston, MA 02110  
Tel: (617) 456-8000  
Fax: (617) 456-8100

**COUNSEL for PLAINTIFFS**

CERTIFICATE OF SERVICE

I hereby certify that on the 6<sup>th</sup> day of March, 2020, I electronically filed the foregoing with the Clerk of using the CM/ECF system which will send notification of such filing to all counsel of record.

/s/ Robert Gilman

**UNITED STATES DISTRICT COURT  
WESTERN DISTRICT OF TEXAS  
AUSTIN DIVISION**

Intellectual Ventures I LLC and Intellectual Ventures II LLC,  Plaintiffs  v.  VMware, Inc.,  Defendant.	Civil Action No. 1:19-cv-01075-ADA
---	------------------------------------

**PLAINTIFFS' RESPONSIVE CLAIM CONSTRUCTION BRIEF**



## TABLE OF CONTENTS

Introduction.....	1
The '686 and '726 Patents .....	2
1. Disputed Terms in the '686 Patent.....	4
A. Modif[y/ied] [a] resource allocation / modify[ing] [the] computer resources allocated to a virtual server ('686 patent claims 5-7) .....	4
B. “resource unavailable messages” / “denied requests to modify a resource allocation” ('686 patent claims 5-7) .....	6
C. “determination that a virtual server is overloaded” ('686 patent claims 5-7).....	10
D. “virtual server” ('686 patent claims 5-7) .....	12
E. “determining that a second physical host can accommodate the requested modified resource allocation” ('686 patent claims 5-7) .....	14
2. Disputed Terms in the '726 Patent.....	16
A. “resource denials” ('726 patent claims 1, 4-5 & 8) .....	16
B. “quality of service guarantee” ('726 patent claims 1 & 4) .....	17
3. Alleged Means Plus Function Terms for the '726 and '686 Patents .....	18
A. '686 patent claim 7 “component” terms (i.e., clauses 1-3 of Ex. A) .....	18
B. '726 patent claim 1, 4 & 5 (i.e., clauses 4-8 of Ex. A) .....	21
The '752 Patent .....	25
4. Disputed Terms in the '752 Patent.....	25
A. “exhausted” ('752 patent claims 1, 9 and 24) .....	25
B. “consumed” ('752 patent claims 1, 9 and 24).....	28
C. “service” ('752 patent claims 1, 3, 9 and 24).....	29
D. Means-Plus-Function Terms .....	29
The '051 Patent .....	31
5. Disputed Terms in the '051 Patent.....	31
A. “virtual server” (claims 1, 3, and 6).....	31
B. “physical interface[s]” (claims 1 and 3).....	33
C. “storing . . .” / “storing . . .” / “receiving . . .” / “determining . . .” / “determining . . .” / “sending . . .” / “using . . .” (claims 1 and 3) .....	34
D. “customer forwarding [table(s)/information]” (claims 1 and 3) .....	37
6. Disputed Terms in the '818 Patent.....	38
A. “hierarchical token bucket resource allocation” / “token(s)” (claims 1, 17, 30, 32 and 42) .....	38

B. “enforcing . . .” / “receiv[e/ing] . . .” / “classify[ing] . . .” / “compar[e/ing] . . .” / “forward[ing] . . .” / “buffer[ing] . . .” (claims 1, 17, 30, 32, 33, 37-39, 42) .....	40
C. “maintain[ing] a connection over a network fabric” (claims 1, 17, 30, 32, 42) ....	42
D. “virtual [network/storage network] interface layer of an application server” (claims 1, 17, 32, 42) .....	44
E. Alleged Means Plus Function Elements (claim 17) (individually set forth in Ex. C) .....	46
Conclusion .....	48

## TABLE OF AUTHORITIES

### Cases

<i>Accent Packaging, Inc. v. Leggett &amp; Platt, Inc.</i> , 707 F.3d 1318 (Fed. Cir. 2013).....	36, 45
<i>ActiveVideo Networks Inc. v. Verizon Communications Inc.</i> , 694 F.3d 1312 (Fed. Cir. 2012).....	1
<i>ACTV, Inc. v. Walt Disney Co.</i> , 346 F.3d 1082 (Fed. Cir. 2003).....	35
<i>Arlington Indus., Inc. v. Bridgeport Fittings, Inc.</i> , 632 F.3d 1246 (Fed. Cir. 2011).....	36, 41, 43
<i>Baldwin Graphic Sys., Inc. v. Siebert, Inc.</i> , 512 F.3d 1338 (Fed. Cir. 2008).....	35
<i>Bell &amp; Howell Document Mgmt. Prods. Co. v. Altek Sys.</i> , 132 F.3d 701 (Fed. Cir. 1997).....	38
<i>Comark Commc'ns, Inc. v. Harris Corp.</i> , 156 F.3d 1182 (Fed. Cir. 1998).....	11, 14, 17, 29
<i>Cox Commc'ns, Inc. v. Sprint Commc'n Co. LP</i> , 838 F.3d 1224 (Fed. Cir. 2016).....	42
<i>Duncan Parking Techs, Inc. v. IPS Grp., Inc.</i> , 914 F.3d 1347 (Fed. Cir. 2019).....	17
<i>Fisher-Rosemount Sys. v. ABB Ltd.</i> , Case No. 4:18-cv-00178, 2019 WL 6830806 (S.D. Tex. Dec. 12, 2019).....	18, 46
<i>Georgetown Rail Equip. Co. v. Holland L.P.</i> , 867 F.3d 1229 (Fed. Cir. 2017).....	41
<i>Graphon Corp. v. Autotrader.com, Inc.</i> , Case No. 2:05-cv-530 (TJW), 2007 WL 1870622 (E.D. Tex. June 28, 2007) .....	15
<i>iFly Holdings LLC v. Indoor Skydiving Germany GmbH</i> , Case No. 2:14-cv-1080-JRG-RSP, 2015 WL 9258264 (E.D. Tex. Dec. 18, 2015).....	13
<i>Inline Plastics Corp. v. EasyPak, LLC</i> , 799 F.3d 1364 (Fed. Cir. 2015).....	6
<i>Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.</i> , 381 F.3d 1111 (Fed. Cir. 2004).....	40
<i>Liebel-Flarsheim Co. v. Medrad, Inc.</i> , 358 F.3d 898 (Fed. Cir. 2004).....	10
<i>Medrad, Inc. v. MRI Devices Corp.</i> , 401 F.3d 1313 (Fed. Cir. 2005).....	10, 33
<i>Medtronic, Inc. v. Guidant Corp.</i> , 465 F.3d 1360 (Fed. Cir. 2006).....	35

<i>O2 Micro Int’l Ltd. v. Beyond Innovation Tech. Co.</i> , 521 F.3d 1351 (Fed. Cir. 2008).....	1, 42
<i>Paragon Sols., LLC v. Timex Corp.</i> , 566 F.3d 1075 (Fed. Cir. 2009).....	41
<i>Renishaw PLC v. Marposs Societa' per Azioni</i> , 158 F.3d 1243 (Fed. Cir. 1998).....	33, 46
<i>Sound View Innovations, LLC v. Facebook, Inc.</i> , Case No. 16-cv-116 (RGA), 2017 WL 2221177 (D. Del. May 19, 2017).....	22
<i>Telcordia Tech., Inc. v. Cisco Sys., Inc.</i> , 612 F.3d 1365 (Fed. Cir. 2010).....	19, 47
<i>Thorner v. Sony Computer Entm’t Am. LLC</i> , 669 F.3d 1362 (Fed. Cir. 2012).....	10, 33, 43
<i>VLSI Tech. LLC, v. Intel Corp.</i> , Case No. 1:19-CV-00977-ADA.....	30
<i>Zeroclick, LLC v. Apple Inc.</i> , 891 F.3d 1003 (Fed. Cir. 2018).....	46
<b>Other Authorities</b>	
MPEP § 2173.05(e) (8 <sup>th</sup> ed. Rev. 2 May 2004).....	15

## Introduction

VMware reveals its claim construction strategy with the first line of its Opening Claim Construction Brief (“VMware Br.”): “This case is immensely complicated.” *VMware Br.* at 1. In fact, VMware has made the claim construction process unnecessarily complicated in an effort to create non-infringement arguments. VMware reads limitations into the asserted claims and conjures up invalidity positions with a kitchen-sink approach to claim construction. VMware employs a number of tactics to further these efforts. VMware’s proposed constructions (1) mix and match terms within claims and even across patents; (2) seek to construe entire claim elements, as well as multiple elements, as single proposed “terms;” (3) argue entire groups of terms should be construed as means-plus-function despite the terms not including the word “means;” and (4) utilize a fifty-three (53) page expert declaration to skew claim language. Below, and in its Opening Claim Construction Brief (“IV Br.”), IV demonstrates that by following well-established claim construction canons, each of the claims and their constituent terms can be well understood by examining the intrinsic record.

VMware misinterprets Federal Circuit law to repeatedly argue that the Court should adopt VMware’s constructions for terms that should simply be construed as plain and ordinary meaning. The case relied on by VMware, *O2 Micro*, stands for the proposition that a term should not be construed with its ordinary meaning if there are multiple possible ordinary meanings or if plain and ordinary will not resolve the parties’ dispute. *O2 Micro Int’l Ltd. v. Beyond Innovation Tech. Co.*, 521 F.3d 1351, 1361 (Fed. Cir. 2008). That is not the case here. IV’s plain and ordinary proposals resolve the claim construction dispute with VMware and are consistent with the intrinsic record and how a person of ordinary skill in the art (POSITA) would have understood them at the time of the invention. *See ActiveVideo Networks Inc. v. Verizon Commc’ns Inc.*, 694 F.3d 1312, 1325-26 (Fed. Cir. 2012) (Federal Circuit rejects argument under *O2 Micro* that a term should not have been construed with its plain and ordinary meaning because doing so did, in fact, resolve the claim construction dispute between the parties).

The patents-in-suit utilize language readily understood by one of ordinary skill in the art. Examining each term individually in the context of the claim in which it appears bears this out.

### **The '686 and '726 Patents**

VMware's efforts to overly complicate this case begin with its approach to United States Patent Nos. RE 44,686 ("the '686 patent") and RE 42,726 ("the '726 patent"). For example, despite the term "quality of service guarantee" appearing in only select claims of the '726 patent, VMware attempts to read it into the construction of nearly all claims across both the '686 and '726 patents. In another example, it points to use of the term "virtual server" in United States Patent No. RE 43,051 ("the '051 patent") to support its construction of the same term in the '686 and '726 patents, but then argues that a different construction of "virtual server" is required in the '051 patent itself.

To help clarify the following arguments with respect to the '686 and '726 patents a short background is instructive. The '686 patent and '726 patent are related and share a specification. They have the same inventors, and both originated from Ensim Corporation. The '726 patent is a reissue of U.S. Pat. No. 6,985,937 and the '686 patent is a continuation of the application that would issue as the '726 patent. Despite sharing a specification, there are key differences in the claims that should be considered. For instance, the '726 patent includes claims with the term "quality of service guarantee" whereas the '686 patent does not include that term at all. VMware glosses over this important distinction. Another difference is that the '726 patent's claims are written in a way that focuses more on components, e.g., 'a virtual server resource monitor,' 'a virtual server resource modifier,' 'a load balancing module,' and a 'dynamic virtual server mover,' and how those components operate together. The '686 patent's claims on the other hand, are written to focus more on the messages and indications that are used by the components to come to the desired result. In some cases, this leads to different terms having similar meanings, like 'resource denials' from the '726 patent and 'denied requests to modify a resource allocation' from the '686 patent. Such a result is not inconsistent, however, because it is based on the specific language and sequence of the claims at issue.

A short background of the '051 patent-in-suit is also helpful in exposing the needless complication VMware creates by cross-referencing the '051 patent in arguments regarding the '686 and '726 patents. While all three patents originated with Ensim Corporation, they are unrelated and share no common inventors. In addition, the inventions taught by the '686 and '726 are quite different than that taught by the '051 patent.

The core concept embodied in the '686 and '726 patents concerns the dynamic modification of resource allocations made to a virtual server by moving the virtual server from one physical host (a server) to another physical host if the virtual server is not getting enough resources to function properly on its current host. Through a series of components, indications and messages, the inventors of the '726 and '686 patents conceived of a way to allow the physical resources used by a virtual server to be dynamically changed if its physical host was overloaded and could not allocate additional resources to that virtual server. That way is to dynamically transfer that virtual server to another host. Because re-distributing such resources on a dynamic basis requires coordination and cooperation among many different components, the sequence of the claimed elements is of great importance in these patents.

The '051 patent, while also dealing with virtualized servers, utilizes virtualization to a different end. The '051 patent's specific combination of virtualized components, tunnels, and physical interfaces allow data centers to be siloed on a per-tenant basis so that each tenant may use the same physical hardware as another tenant without risk of security breach or hardware failure. Furthermore, because the tenant's servers are virtualized, the network addresses of one tenant's virtual servers may overlap with the network addresses of another tenants, virtual server, which further lessens the risk of a security breach while providing each tenant with a lot of IP addresses.

It is against these distinct backdrops that the claims of these three patents must be analyzed. When examined in the proper context, it becomes clear that IV's proposed constructions are appropriate, and that VMware is attempting to needlessly sow confusion.

## 1. Disputed Terms in the '686 Patent

### A. **Modif[y/ied] [a] resource allocation / modify[ing] [the] computer resources allocated to a virtual server ('686 patent claims 5-7)**

IV's Proposed Construction	VMware's Proposed Construction
"modif[y/ied] set of functions and features of a physical host(s) used in implementing tasks for each virtual server" / "modify[ing] a set of the functions and features of a physical host(s) used in implementing tasks for each virtual server"	"modif[y/ied] [a] quality of service guarantee" / "modify[ing] [the] quality of service guarantee of a virtual server"

In its Opening Brief, VMware mischaracterizes the intrinsic record and IV's proposed construction in an attempt to read unsupported limitations into the disputed terms. First, VMware mischaracterizes the prosecution history of the '686 patent by stating that during prosecution of the '937 patent (parent of the '686 patent), the applicant disclaimed any meaning of "resource allocation" other than "quality of service guarantee." *See* VMware Br. at 3. Second, VMware tries to support this "quality of service guarantee" construction with support from one of the claims of the '726 patent, despite the fact that this phrase is not found in any of the claims of the '686 patent. Third, in arguing that IV's construction reads out "allocation" and seeks to construe only "modifying a resource" VMware is in effect disregarding the wording of the claim. *See id.*

VMware's argument that there was an "unmistakable prosecution history disclaimer" requires a misreading of the '686 patent's prosecution history. The applicant's statement, relied on by VMware, that "a resource allocation for a process is specified as a quality of service guarantee," referred to the specific claim language of the embodiment captured by then claim 1, not to the scope of the invention as a whole (and certainly not to claims entirely devoid of the term "quality of service guarantee").

Claim 1, as written at the time of the applicants' statement, is reproduced below:

1. In a computer system including one or more physical hosts, a method for dynamically modifying the computer resources allocated to a process, the method comprising:

**determining that a resource allocated to a process on a first physical host is overloaded; and**

**in response to the resource being overloaded, increasing a quality of service guarantee for the process.**



(emphasis added). As can be seen by viewing the prosecution history, the claim then being prosecuted required that the resource allocation be specified as the claimed quality of service guarantee.

A closer look at the evolution of the claim during the prosecution of the '937 patent confirms this. For example, in the March 15, 2004 Response to Office Action, the applicant amended claim 1 to further clarify this point, adding “*the portion of the resources being specified as a quality of service guarantee*” and “in response to *the portion of the resources* allocated to the process being overloaded, increasing the quality of service guarantee for the process.” See Second Declaration of Jonathan R. DeBlois (hereinafter “2d DeBlois Decl.”)<sup>1</sup> Ex. R at 2. Giving key context to the amendment, the applicant stated that it was made because the examiner and applicant disagreed as to whether a resource could both meet a quality of service guarantee and be overloaded, as well as to clarify the term “quality of service guarantee.” *Id.* at 14.

The examiner, however, was not convinced, and in the April 21, 2005 Response to Office Action, the applicant canceled all claims that included “quality of service guarantee” being modified or increased as a result of an overload. Instead, the applicant relegated the phrase to the preamble of what would issue as claim 1, noting only that “resources allocated to the virtual server being specified as a quality of service guarantee.” In the body of the claim, in place of the language previously claiming an increase or modification of the “quality of service guarantee” to cure an overload, the applicant added the disputed term “modify a resource allocation for the virtual server.” *Id.* at 19-22. The examiner subsequently allowed the claims.

As can be seen by review of the entire prosecution history there was no disclaimer as to the resource allocation terms because the applicant’s arguments were directed to the scope of “quality of service guarantee.” Indeed, in every instance where the applicant made statements regarding “quality of service guarantee” the comments were focused on qualifying the phrase as specifying or representing a dynamic resource allocation. This makes logical sense, since as the

---

<sup>1</sup> All exhibit citations are to the Declaration of Jonathan R. DeBlois attached to IV’s Opening Claim Construction Brief unless otherwise noted.

examiner correctly noted, the term quality of service guarantee typically refers to a static allocation, not an allocation that can be dynamically modified. *See* VMware Br. at 4.

Furthermore, the fact that the applicant used the term “quality of service guarantee” in only three of the twenty-five claims of the ’726 patent supports IV’s proposed construction, not VMware’s argument. For instance—even setting aside the fact that the phrase is not used in any claims of the ’686 patent for the moment—its use in the preamble of claim 1 of the ’726 patent is simply to note that in the claimed system “the computer resources allocated to the virtual server [are] specified as a quality of service guarantee.” *See* Ex. F at claim 1. This is not saying that the resource allocations are quality of service guarantees, rather, one way in which a resource allocation can be represented is through a quality of service guarantee. This point is further highlighted by the applicant choosing to include the phrase in certain claims of the ’726 patent, while leaving it out of others, (and the ’686 patent), entirely. *See Inline Plastics Corp. v. EasyPak, LLC*, 799 F.3d 1364, 1371 (Fed. Cir. 2015) (error to construe disputed term to include an element expressly claimed in subsequent claim but not the claim at issue).

Finally, VMware’s argument that IV’s proposed construction reads out the term “allocation” is unfounded. As one of ordinary skill in the art would readily understand, for resources to be used by a virtual server those resources necessarily must be allocated. Therefore, IV’s proposed construction fully accounts for the disputed ‘allocation’ term. That being said, if the court is so inclined, IV would be amenable to removing the iterations of “each” in its proposed construction and replacing them with “the” to allay any potential concerns regarding its construction being unclear.<sup>2</sup>

**B. “resource unavailable messages” / “denied requests to modify a resource allocation” (’686 patent claims 5-7)**

IV’s Proposed Construction	VMware’s Proposed Construction
“an indication that a request by the virtual server cannot be immediately serviced” / “a request by	“indications that requests by the virtual server for additional resources are either implicitly or

<sup>2</sup> This would make IV’s construction “modif[y/ied] set of functions and features of a physical host used in implementing tasks for the virtual server.”

IV's Proposed Construction	VMware's Proposed Construction
the virtual server that cannot be immediately serviced"	explicitly denied, resulting from requests to modify a resource allocation" <sup>3</sup> see also construction of "modify a resource allocation"

VMware hinges much of its argument with respect to this term on its incorrect statement that "[b]oth parties have proposed this term for construction." VMware Br. at 7. In actuality, VMware proposed the term "resource unavailable messages resulting from denied requests to modify a resource allocation" as a single term for construction, while IV proposed the two constituent terms "resource unavailable messages" and "denied requests to modify a resource allocation." Clarifying which terms were actually proposed by which party renders VMware's argument that IV's construction attempts to read out the phrase "resulting from denied requests to modify a resource allocation" incorrect.

IV did not propose the entire phrase "resource unavailable messages resulting from denied requests to modify a resource allocation" for construction because the term 'resulting from' requires no construction. Thus, the dispute lies with the two terms connected by the 'resulting from' language, i.e. "resource unavailable messages" and "denied requests to modify a resource allocation" (each of which IV proposed to construe). With those two terms construed, it is clear that when they are connected using the phrase 'resulting from' that one is the result of the other. VMware itself acknowledges this in its proposed construction, which simply construes "resource unavailability messages" and then tacks on the remainder of the disputed term verbatim. *See* VMware Br. at 7 ("indications that requests by the virtual server for additional resources are either implicitly or explicitly denied, *resulting from denied requests to modify a resource allocation*") (emphasis added).

Neither is IV's construction "circular" as VMware alleges. This becomes particularly evident when the disputed terms are read in the proper order, i.e., the order in which they occur within the computer system. For illustrative purposes the graphic below uses the relevant

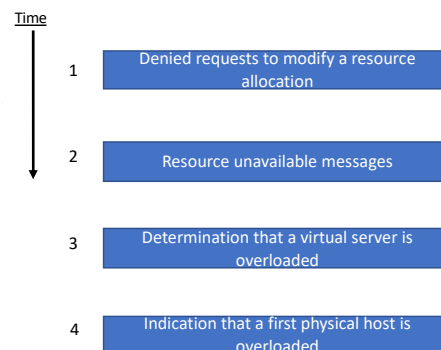
---

<sup>3</sup> VMware proposed this construction as a 'compromise' during the parties' efforts to reduce the number of terms but did not indicate to IV that it would adopt that proposal as its new construction. As a result, IV's Opening Brief included an older version of VMware's proposed construction.

elements of claim 5 of the '686 patent as an example to show how the order of the elements is key to construing the disputed terms:

5. A method performed by a computing device, having a processor and memory, for modifying the computer resources allocated to a virtual server operating in a first physical host of multiple physical hosts, comprising:

receiving <sup>4</sup>an indication that a first physical host is overloaded, wherein the indication is based on a determination that a virtual server is overloaded and wherein the <sup>3</sup>determination that a virtual server is overloaded is based on one or more <sup>2</sup>resource unavailable messages resulting from denied requests to modify a resource allocation; <sup>1</sup>



As can be seen above, first there is a *denied request to modify a resource allocation*.

Subsequently, as a result of that denied request, there is a *resource unavailable message* generated. Based on the resource unavailable message, a *determination that a virtual server is overloaded* is then made. Finally, after it is determined that the virtual server is overloaded, an *indication that a first physical host is overloaded* is received. Importantly for the disputed terms here, the ‘resource unavailable messages’ are simply the result of the ‘denied request to modify a resource allocation,’ which are then used to make ‘the determination that a virtual server is overloaded.’

With this proper sequence in mind, the specification of the '686 patent describes how the two disputed terms should be construed. Because the claim language requires the ‘*determination that a virtual server is overloaded*’ be based on ‘*resource unavailable messages*,’ which themselves are representative of ‘*denied requests to modify a resource allocation*,’ the construction analysis begins with how the specification describes the ‘*determination that a virtual server is overloaded*.’ At column 2 lines 62-65 such a determination is disclosed as being based on the monitoring of ‘resource denials.’ The specification further defines the term ‘resource denials’ as “any request by the virtual server that cannot be immediately serviced.” Ex. E at 65-66. Therefore, it logically follows that the ‘*denied requests to modify a resource*

*allocation*’ equate to ‘resource denials,’ and should be properly construed as “a request by the virtual server that cannot be immediately serviced.” Following this same logic, the ‘*resource unavailable messages*’ should be construed as “an indication that a request by the virtual server cannot be immediately serviced,” since the indication is how the denials are quantitatively accounted for by the computer system.

This makes perfect sense when considering the operation of the claimed invention in its entirety and as reflected in the specification. The invention is a computer system that can only take action based on computer instruction sets. A denied request to modify a resource allocation, alone is not enough to cause the computer system to take action. Instead, in order for the system to react to the denied request, an instruction must cause the system to do so. Since in the claim, the denied requests must be quantitatively accounted for to make a determination of a virtual server overload, the computer is programmed to send a message that a request to modify a resource allocation has been denied (i.e., the claimed ‘resource unavailable messages’). *See* Ex. E at 7:47-55. These indications then become representative of the denied requests and are used as a quantitative measurement to determine whether the virtual server is overloaded. *Id.* Accordingly, IV’s proposed construction is consistent with both the intrinsic record and the actual purpose and operation of the claimed inventions. *See* Ex. E at column 8.

In addition, despite VMware’s conclusory argument to the contrary, IV’s proposed constructions of the disputed terms are entirely consistent with its proposed construction of “modified resource allocation.” Common sense dictates that a ‘denied request to modify a resource allocation’ is quite different than a ‘modified resource allocation.’ A modified resource allocation is a resource allocation that has been modified, i.e., something tangible; a new resource allocation. A denied request to modify a resource allocation, on the other hand, is a denied request, i.e. a resource denial. There is no inconsistency between IV’s proposed construction of the disputed terms and its construction of ‘modified resource allocation.’ The terms are two distinct concepts that by their very nature cannot be distilled into a single construction.

**C. “determination that a virtual server is overloaded” (‘686 patent claims 5-7)**

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	“determination that an average number of resource denials for a virtual server is beyond a pre-configured threshold” See also construction of “resource denials”

VMware’s arguments with respect to this term fail to take into account the knowledge of one of skill in the art and the full intrinsic record. First, VMware claims that because the parties agreed on the construction of “indication that a first physical host is overloaded” the disputed term must be construed as something other than plain and ordinary meaning, since it is similarly a “technical term.” *See* VMware Br. at 9. The fact that the disputed term is “technical,” however, is ultimately irrelevant as to whether it should be construed as having its plain and ordinary meaning. The litmus test for claim construction is not whether a term is technical, but rather, whether it requires a specific construction because the patentee either made a clear disclaimer or acted as his own lexicographer. *See Thorner v. Sony Computer Entm’t Am. LLC*, 669 F.3d 1362, 1365 (Fed. Cir. 2012). Otherwise, the claim should be given its plain and ordinary meaning as understood by one of skill in the art at the time of the invention. *See id.* In this case, one of skill in the art would readily understand that the disputed term is being used in its well-understood and customary manner. *See* IV Br. at 18-20; *see also Liebel-Flarsheim Co. v. Medrad, Inc.*, 358 F.3d 898, 913 (Fed. Cir. 2004).

VMware’s next point that the disputed term and the ‘physical host overloaded’ term are different, and thus require different constructions is (i) irrelevant and (ii) in any event, consistent with IV’s position. Namely, that the term an ‘indication that a first physical host is overloaded,’ in light of the specification and claim elements, requires a construction other than plain and ordinary meaning, while “determination that a virtual server is overloaded does not. Turning to the construction at hand, a virtual server is a well-known and commonly used term among those of skill in the art and means exactly what it says; a virtualized server. *See* Declaration of Dr. Robert Akl, D.Sc. (hereinafter “Akl Decl.”) at ¶ 31. An overload in the context of a virtual server is equally well understood and is further explained by the language of the claim itself,

namely, that it is based on resource unavailability messages generated from denied requests to modify a resource allocation. *See id.* As described in detail in the previous section, the patent's specification further discloses exactly what is meant by resource denial and unavailability messages. *See supra* Section 1(B). Therefore, the disputed term has a well-understood plain and ordinary meaning and it should be construed as the same.

VMware's argument that the specification provides a "clear definition" of this term at column 5 line 29 is incorrect. Indeed, far from reciting a clear definition, column 5 line 29 describes "a flowchart of an embodiment" of this portion of the invention. Ex. E at 5:29. Reading a limitation into the claims from a preferred embodiment violates the most fundamental claim construction tenets. *See Comark Commc'ns, Inc. v. Harris Corp.*, 156 F.3d 1182, 1186-87 (Fed. Cir. 1998). Moreover, this portion of the specification does not describe how a virtual server is determined to be overloaded. Rather, it is describing how it is determined whether or not a virtual server resource is overloaded. Specifically, VMware's purported support for its construction is discussing "whether a particular virtual server resource is overloaded," and states that to make a determination of the same it uses "a number of well-known techniques," finally noting that the determination is with respect to whether the virtual server is "overloaded for the corresponding resource." Ex. E at 5:42-43. The disclosure relied on by VMware, at best, only lends ambiguity to its proposed construction.

VMware alleges that its proposed construction is the only possible option that would not render the disputed term indefinite for lack of written description. VMware, however, fails to point to any evidence supporting this allegation. To the contrary, the intrinsic record is replete with evidence that the disputed term is being used in its plain and ordinary meaning. *See, e.g.*, Ex. E at Abstract, 2:62-65, 5:15-17, 8:5-55. Therefore, the disputed term is not indefinite and should be construed to have its plain and ordinary meaning.<sup>4</sup>

---

<sup>4</sup> VMware's allegation that IV "refused to explain what it believes the plain and ordinary meaning" of the disputed term is or "otherwise explain why it disagrees with VMware's proposal," is simply wrong. IV explained to VMware that its proposed construction was reading limitations into the claim and thus improper. IV also explained, via



**D. “virtual server” (’686 patent claims 5-7)**

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning, in the alternative “a virtual machine that resides on a physical server and uses the physical server’s resources but has the appearance of being a separate dedicated machine”	“a process executing on a host computer that accepts communication requests, and that is capable of receiving a quality of service guarantee from a physical host”

VMware attempts to limit the common term in the art “virtual server” in two ways and neither has support in the intrinsic record. First, VMware tries to limit the patent’s virtual server to “a process executing on a host computer that accepts communication requests” by principally relying on a contorted reading of the unrelated ’051 patent. Second, VMware asks the Court to read a limitation about what the virtual server can do from a description of the preferred embodiments into the claim. That violates two claim construction canons.

IV’s proposal of plain and ordinary meaning should be adopted because the disputed term is a basic building block of virtual computing and has a well-known and understood meaning among those with skill in the art. *See* Akl Decl. at ¶ 31. Furthermore, nothing in the intrinsic record of any of the asserted patents rises to the level of a clear disclaimer of scope or indicates that the patentee acted as his own lexicographer. For example, the ’686 specification goes into great detail regarding the background of the invention and the field of technology, all of which supports IV’s proposed construction of plain and ordinary meaning. In doing so the patentee discusses the concept of virtual servers replacing physical servers which were owned by an ISP and leased in their entirety to a single customer. The patentee notes that many customers, however, do not want (or need) the processing power of an entire physical server. A better solution would be to use a single physical server to host multiple virtual servers that customers could customize to their needs. This clearly indicates that the patentee is using the term in its plain and ordinary meaning, i.e., a server that is virtualized. *See generally*, Ex. E at column 1; *see also iFly Holdings LLC v. Indoor Skydiving Germany GmbH*, Case No. 2:14-cv-1080-JRG-

---

phone and email conversations, that the term should be given its plain and ordinary meaning as understood by one of skill in the art because the intrinsic record shows this to be the case, and that by its very nature a plain and ordinary construction requires no further definition.



RSP, 2015 WL 9258264, at \*19 (E.D. Tex. Dec. 18, 2015) (“in other words, the plain and ordinary meaning of the phrase ‘said chamber having a width’ is exactly what it plainly states, the chamber has a width.”).

On the other hand, the intrinsic record provides no support to limit “virtual server” with “a process executing on a host computer that accepts communication requests . . . .” The claims themselves and the intrinsic record surrounding the term contradict VMware’s proposed limitation. First of all, VMware fails to note that the prosecution history of the ’937 patent demonstrates that the disputed term was not meant to be limited to or defined as a “process.” Specifically, up until the October 5, 2004 Response to Office Action the claims included the term “process” and not “virtual server.” *See* 2d DeBlois Decl., Ex. R at 23-40. In the aforementioned office action response, however, **the applicant amended the claims to remove “process” and instead use “virtual server.”** *Id.* Had the applicant intended the disputed phrase to be limited to a process the claims would not have been amended to remove that exact term in favor of “virtual server.”

This is further supported by the specification of the ’686 patent. For example, the following excerpt is a quote from column 2 and is helpful in exposing VMware’s flawed construction:

Different customers have **different virtual server needs**. For example, a company A **providing large quantities of data and information to its employees and customers will want to ensure that its virtual servers are always available** to perform a large number of tasks. Company A may be willing to pay a premium for a guaranteed high quality of service, **with high server availability and large amounts of processing power** always on-call.

Ex. E at 2:1-8. The patentee is describing a virtual server just as anyone with skill in the art would expect and understand, not in terms of something that differs from or alters its plain and ordinary meaning. A customer needs a server to provide data and information to employees and customers; the server must always be available and have sufficient processing power to perform a large number of tasks. The server in this case simply happens to be virtual. What it is not, on the other hand, ‘is a process executing on a host computer’ as VMware suggests. A process

cannot have “high server availability” or “large amounts of processing power.” A process also cannot “provide large quantities of data and information.” Perhaps most importantly, there is nothing about a process that is virtual, in other words, every computer can (and does) run processes to implement all sorts of functions (including implementing a virtual server). VMware’s construction interjects confusion and uncertainty into the claims and ignores that virtual server was specifically used in place of ‘process.’<sup>5</sup>

VMware’s second limitation, requiring a virtual server to be “capable of receiving a quality of service guarantee from a physical host,” lacks intrinsic support and is inconsistent within the context of the claims. VMware’s proposed limitation comes from a portion of the specification clearly labeled as describing embodiments. Ex. E at 49-51. VMware cites to one embodiment where virtual servers are capable of receiving quality of service guarantees and claims that the construction therefore must include that phrase. But VMware fails to note that immediately preceding that description the patentee states that “[r]eference will now be made in detail to several embodiments of the present inventions,” language clearly indicating that the invention is not limited to those embodiments. *Id.* Additionally, as discussed above *supra* Section 1(A) there are claims of the asserted patents that include the term ‘quality of service guarantee’ and claims that do not. Reading a limitation from a specific embodiment violates the most basic of claim construction principles. *See Comark Commc’ns, Inc.*, 156 F.3d at 1186-87.

**E. “determining that a second physical host can accommodate the requested modified resource allocation” (’686 patent claims 5-7)**

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning; in the alternative: “determining that a second physical host can accommodate the requests by the virtual server that could not be immediately serviced”	Indefinite, or in the alternative “determining that a second physical host can accommodate the denied request to modify a resource allocation”

The disputed term is not indefinite because it clearly refers back to the prior iteration of “denied requests to modify a resource allocation,” and therefore, has sufficient antecedent basis.

<sup>5</sup> While IV’s alternate proposal more accurately captures the scope of the claims and spirit of the invention and should be adopted rather than VMware’s erroneous construction, however, IV’s position remains that a construction of plain and ordinary meaning is most appropriate here.

The requirement of antecedent basis is a rule of patent drafting. The MPEP states that a failure to include explicit antecedent basis is not always fatal to a claim. MPEP § 2173.05(e) (8<sup>th</sup> ed. Rev. 2 May 2004). The Federal Circuit has clarified that lack of antecedent basis does not render a claim indefinite “if the scope of a claim would be readily ascertainable by those of skill in the art.” *See Graphon Corp. v. Autotrader.com, Inc.*, Case No. 2:05-cv-530 (TJW), 2007 WL 1870622, at \*11 (E.D. Tex. June 28, 2007) (citation omitted). When the instant claim is read in its entirety and in light of the specification the scope of the disputed term is readily apparent to one of skill in the art. For instance, the claim element immediately preceding the disputed term recites<sup>6</sup>, in relevant part, “resource unavailable messages resulting from denied requests to modify a resource allocation.” Furthermore, the specification clearly discloses what a denied request to modify a resource allocation is, how to monitor for it, and gives at least one embodiment complete with exemplary algorithms. *See, e.g.*, Ex. E at 11:12-57. One of skill in the art would clearly understand that the disputed term is referring back to the denied request to modify a resource allocation and therefore would also understand that the claimed second physical host would be required to have sufficient resources to satisfy the previously requested and denied resource allocation. *See* Akl Decl. at ¶ 32.

In addition, the same evidence supports IV’s position that the disputed term should be given its plain and ordinary meaning as opposed to VMware’s alternate proposed construction, which convolutes the term and, therefore, would be unhelpful to the jury. As discussed more fully in IV’s Opening Brief at Section B(7), there is a plain meaning in the art with respect to the disputed term’s constituent parts, and therefore, one of ordinary skill in the art would understand exactly what is meant by the disputed term. VMware’s proposal on the other hand, just inserts “the denied request to modify a resource allocation” in place of the disputed terms “the requested modified resource allocation.” It further adds “*see also* construction of “modify a resource allocation.” Not only is this unhelpful since it inserts one disputed term in place of another under

---

<sup>6</sup> Claim 5 of the ‘686 patent is being used here as representative.

the guise trying to “remedy the antecedent basis problem,” it is unclear what VMware’s alternate proposed construction even is in light of the ambiguous “see also” citation it includes. IV’s alternate construction instead incorporates its construction of the term “denied request to modify a resource allocation,” so that its alternate proposal correctly reflects the intrinsic record and provides guidance to the factfinder as to the meaning of the disputed term. *See supra* Section 1(B).

## 2. Disputed Terms in the ’726 Patent<sup>7</sup>

### A. “resource denials” (’726 patent claims 1, 4-5 & 8)

IV’s Proposed Construction	VMware’s Proposed Construction
“indications that requests by the virtual server cannot be immediately serviced”	“indications that requests by the virtual server for additional resources are either implicitly or explicitly denied”

As an initial matter VMware notes that the parties have agreed on the first part of the proposed construction—“indications that requests by the virtual server”—for the disputed term despite the fact that IV’s original proposal did not pluralize “indication.” Prior to VMware’s Opening Brief, IV was unaware that VMware was amenable to agreeing on this portion of the construction since VMware made no mention of it during the meet and confer process or the parties’ follow-ups. Nevertheless, IV agrees to modify its proposed construction to reflect this agreement.

The remainder of VMware’s proposed construction, however, should be rejected because, as discussed in IV’s Opening Brief, Section B(3), it improperly limits the disputed term to one narrow embodiment despite the intrinsic record defining the term more broadly. VMware draws its proposed construction from a cite to the specification that describes “instances wherein a request for additional resources is either implicitly or explicitly denied.” *See* Ex. F at 7:51-54. This cite is in the context of discussing a specific embodiment where the virtual server resource monitor monitors different types of resource denials. *Id.* at 5:41-43. The specification then gives

<sup>7</sup> The ’726 patent terms that overlap with ’686 terms discussed above can be found in Exhibit 7 to VMware’s Opening Brief. These claim terms should be construed as IV proposes for the reasons stated above with respect to the corresponding ’686 patent term.

a description of that embodiment’s ‘resource denials’ which VMware adopts as it proposed construction. While an “explicit” or “implicit” denial can be types of resource denials, it is improper to limit the disputed term to this narrow exemplary embodiment, particularly in light of the broader disclosure that captures the term’s full scope. *See Comark Commc’ns, Inc.*, 156 F.3d at 1186-87. Indeed, IV’s proposed construction is drawn from the patent’s Summary of Invention that provides a broader disclosure of “resource denials.” Ex. F at 2:55-61. As such, VMware is not only asking the Court to read in a limitation, it is also seeking to read out explicitly disclosed claim scope. *See Duncan Parking Techs, Inc. v. IPS Grp., Inc.*, 914 F.3d 1347, 1362 (Fed. Cir. 2019).

**B. “quality of service guarantee” (’726 patent claims 1 & 4)**

IV’s Proposed Construction	VMware’s Proposed Construction
“a guaranteed resource allotment which can be dynamically increased/modified”	“information that specifies a guaranteed amount of an assigned resource, and that can be dynamically increased/modified”

VMware argues that IV’s construction provides no clarity “on what is being guaranteed” despite the plain language of IV’s proposal specifying that an “amount of an assigned resource” is the item that is guaranteed. IV’s construction is consistent with the specification’s disclosure of how resource allocations are linked to “quality of service guarantee.” IV’s proposed construction diverges from the specification only by using the synonym “allotment” for “allocation.” Since “allotment” is synonymous with “allocation”<sup>8</sup> it follows that IV’s proposed construction of “a guaranteed resource allotment which can be dynamically increased/modified” is much true to the specification than VMware’s proposal. The specification illustrates this point in the background section when noting “[h]owever, customers may be unable to anticipate the exact amount of resources they will require, and a static **assignment of a particular resource allocation** limit may not allow the virtual server system to adapt to changing customer needs . . . a better **resource allocation model** . . .” *See* Ex. F at 2:14-19. This passage both links the term

<sup>8</sup> IV would be willing to change its proposed construction to use “allocation” in place of “allotment” to avoid any potential confusion, if the Court is so inclined.

‘assignment’ to something that can’t be altered, and also indicates an ‘assignment’ is distinct from the phrase ‘resource allocation.’

On the other hand, VMware’s proposal of ‘information that specifies a guaranteed amount of an assigned resource’ mixes terms and concepts from other unrelated parts of the patent into “quality of service guarantee.” For instance, VMware’s construction inserts “information that specifies” and “an assigned resource” both of which interject ambiguity into the disputed term in light of the specification’s disclosure. ‘Quality of service guarantee information’ is described in the specification as distinct from a ‘quality of service guarantee.’ See Ex. F at 4:39-60. The former being how a quality of service guarantee is represented in some embodiments, such that the information can be stored in a “parameter table,” while the latter is the actual resource allotment that the virtual server is guaranteed. *Id.* at 2:38-46. Furthermore, the phrase ‘assigned’ appears only twice in the specification both in the context of criticizing the prior art as teaching “fixed” or “static” assignment of resources. See *id.* at 2:5-6, 2:14-16. This contradicts the patented invention’s teaching of “dynamic” resource allocation and VMware’s own proposed construction.

### **3. Alleged Means Plus Function Terms for the ’726 and ’686 Patents**

#### **A. ’686 patent claim 7 “component” terms (i.e., clauses 1-3 of Ex. A)**

As discussed in detail in IV’s Opening Brief, Section B(10)(i), the ‘component’ claims of the ’686 patent (listed in full in Ex. A, clauses 1-3) are not subject to § 112 ¶ 6 because: (1) “means for” is not present in any claims, (2) the prefatory language used does not consist of merely a nonce word followed by function, and (3) the claim language itself provides sufficient structure to avoid the application of § 112 ¶ 6. See *Fisher-Rosemount Sys. v. ABB Ltd.*, Case No. 4:18-cv-00178, 2019 WL 6830806, at \*15-\*16 (S.D. Tex. Dec. 12, 2019). IV will not re-hash these arguments herein in the interest of brevity. In its Opening Brief, however, with respect to Clauses 1-2, VMware argues that the ’686 specification fails to disclose any algorithm that performs the claimed functions. Assuming that the terms are found to be subject to § 112 ¶ 6—

which they should not be—the specification does in fact disclose corresponding structure and algorithm for these disputed terms.

Clause 1, which claims the function of “receiv[ing] an indication that a first physical host is overloaded,”<sup>9</sup> has sufficiently disclosed structure in the intrinsic record, including within the claim itself. This fact would be easily understood by one of skill in the art, which is the standard for the sufficiency of structure in a § 112 ¶ 6 analysis. See *Telcordia Tech., Inc. v. Cisco Sys., Inc.*, 612 F.3d 1365, 1377 (Fed. Cir. 2010) (“the specification need only disclose adequate defining structure to render the bounds of the claim understandable to an ordinary artisan.”).

First, the claim elements preceding the “component” element at issue here provide structure for this function, specifying that the ‘component configured to receive’ is within a physical host that includes one or more processors and one or more memories. These are structural terms that in combination with further description in the specification—as discussed below—clearly disclose the corresponding structure. The ’686 patent additionally discloses structure corresponding to the claimed function in Figure 1, 2A, 2B, and 5, and the related descriptions of the same. Specifically, Figure 1 shows that the invention includes a “dynamic resource configuration module 100” which comprises, among other things, a “physical host load balancer 130.” Figure 5 shows that within the physical host load balancer 130 there is a “physical host resource monitor 540 and load balancing calculator 530.” The specification discloses that:

physical host load balancing module 130 receives a signal 510 from the virtual server resource modifier 120 indicating that virtual server 162B requires an increased resource allocation. This signal is used as an input 520 into the load-balancing calculator 530. The load-balancing calculator 530 also requests and receives as input the current physical host resource loads 535 from the physical host resource monitor 540. The physical host resource monitor 540 performs periodic physical host resource checks 545 upon the group of physical hosts . . . the load-balancing calculator 530 determines whether physical host 160A is capable of supporting the request for additional virtual server 162B resource 510. If the

---

<sup>9</sup> VMware claims that the entire element should be considered the function, however, the portion of the limitation following IV’s proposed function is merely further description of the claimed “indication,” not the function itself which is “receiving an indication that a first physical host is overloaded.” The presence—or lack thereof— of the specific properties of the indication is an infringement/non-infringement issue, not a claim construction issue.



resource request exceeds the available resources of physical host 160A, the load-balancing calculator 530 determines that physical host 160A is overloaded.

Ex. E at 11:4-26. (emphasis added). Accordingly, the load-balancing calculator 530, which is within the physical host load balancing module 130 contained within the dynamic resource configuration module 100, is the structure that receives the indication that a first physical host is overloaded. The indication itself is the current physical host load on 160A received from the physical host resource monitor 540 compared to the request for additional resources 510 by virtual server 162B. Combined with the flow diagrams in Figure 2A and 2B, as well as the structural limitations of the claim itself, this provides sufficient disclosure to one of skill in the art of what structure performs the claimed function. *See* Akl Decl. at ¶ 33.

Clause 2 finds similar structural support in the claim and throughout the specification for the function of “determin[ing] that a second physical host can accommodate the requested modified resource allocation.” For example, the starting point is again within claim 7 itself which specifies that the function is performed within a physical host server with one or more memories and one or more processors, as discussed above. Figure 1, Figure 5, and the specification portions in which they are discussed provide additional structure, as also discussed above. Furthermore, picking up where the above block quote ended, the specification describes in detail how the physical host load balancer 130 and the load-balancing calculator 530 perform the claimed function:

[T]he load-balancing calculator 530 uses an easiest fit heuristic to find the physical host that has the most available resources. Each different type of resource is associated with an ordinal and a weight . . . [discloses easiest fit heuristic algorithm]<sup>10</sup> . . . [u]sing the easiest fit heuristic, the load-balancing calculator 530 will select the physical host with the largest weighted resource availability to receive the virtual server 162B (in the above example of Fig. 5, physical host 160B) . . . the load balancing calculator 530 sends 550 a signal 560 to the dynamic virtual server mover 140 indicating that the virtual server 162B is to be transferred to physical host 160B.

Ex. E at 11:27-57 (emphasis added). Therefore, the structure corresponding to the claimed function is set out in detail and includes, the load-balancing calculator 530, which is within the

---

<sup>10</sup> For the specific ‘easiest fit heuristic’ formula *see* ‘686 Specification at 11:30-45.



physical host load balancing module 130 contained within the dynamic resource configuration module 100. When viewed in light of the easiest fit heuristic formula, the flow diagrams in Figure 2A and 2B, as well as the structural limitations of the claim itself, this provides sufficient disclosure to one of skill in the art of what structure performs the claimed function. *See* Dr. Akl Decl. at ¶ 34.

With respect to Clause 3, IV agrees with VMware's identified function and a portion of the corresponding structure that it has identified, however, VMware's structure is too narrow and leaves out key structural elements that play a part in performing the claimed function. For example, VMware argues that virtual server mover 140 as shown in Figure 6 is the entirety of the corresponding structure for the function. VMware's structure proposal, however, fails to account for the 'generate a physical host transfer signal' portion of the claimed function. The virtual server mover 140 only performs the transfer of the virtual machine from a first to a second host, it does not perform the generating of the transfer signal. *See* Ex. E at 11:27-57.

Therefore, at least the load balancing module 130 and load balancing calculator 530, which are components of dynamic resource configuration module 100, perform the 'generate' function. *See id.* at 11:54-57; Fig. 5. Furthermore, in order to generate the transfer signal, the physical host load balancer 130 needs to receive information from the virtual server resource monitor 120 relating to virtual server 162B's resource request, as well as information from the physical host resource monitor 540 relating to the resource loads on all physical hosts 160A-C. All of these components are located within dynamic resource configuration module 100, and should all be included in the corresponding structure. *See* Akl Decl. at ¶ 35.

**B. '726 patent claim 1, 4 & 5 (i.e., clauses 4-8 of Ex. A)**

As an initial matter, the claim elements reflected in Clauses 4-8 of Ex. A should not be subject to § 112 ¶ 6 for the reasons set forth in IV's Opening Brief, Section 10(ii)-(vi), namely, because one of ordinary skill in the art would be well-informed of the structure of the elements from the claim language and specification disclosures. *See* Akl Decl. at ¶ 36. Should the Court

disagree and find that § 112 ¶ 6 does apply, however, each disputed element has sufficient corresponding structure identified either in the claim itself, specification, or both. *Id.*

Additionally, VMware attempts to read the term “creating” into each one of its proposed functions, despite the fact that only claim 4 actually claims “program code for creating,” while claims 1 and 5 make no mention of the term. VMware should not be allowed to limit the function of an alleged means-plus-function element to a function—“creating”—that is not found in the disputed term. Therefore, IV will address the iterations of the disputed terms of claims 1 and 5 herein. As for the disputed “program code” elements of claim 4, IV agrees with VMware’s proposed function to the extent it contains “creating,” but disagrees with VMware’s proposed structure for the same reasons discussed above with respect to the non-program code elements.

Finally, as with the alleged means-plus-function elements in the ‘686 patent discussed above, VMware seems to misapprehend the scope of what is considered an “algorithm.” For purposes of § 112 ¶ 6, when an algorithm is required to support a function (which it is not here), that algorithm does not have to take the form of a mathematical formula. Anything adequate to permit one of ordinary skill in the art to perceive the bounds of the invention is sufficient, and can be presented in the specification read as a whole in any understandable terms including as a mathematical formula, in prose, or as a flow chart, or in any other manner that provides sufficient structure. *See Intelligent Automation Design, LLC*, 2020 WL 486830, at \*3-\*5.

i. “virtual server resource monitor” term, clause 4 of Ex. A

The term ‘virtual server resource monitor’ as claimed in claims 1 and 5 is required to, (a) monitor resource denials, and (b) send a virtual server overload signal in response. The claim itself provides that this takes place in a physical host server including a processor and memory. Contrary to VMware’s conclusory assertions, ‘server’ can provide structure and has a well-known meaning in the art. *See Sound View Innovations, LLC v. Facebook, Inc.*, Case No. 16-cv-116 (RGA), 2017 WL 2221177, at \*10-\*11 (D. Del. May 19, 2017). Furthermore, the specification and figures provide additional structure beyond simply just virtual server resource

monitor 110, as proposed by VMware. The summary of the invention notes that “a dynamic resource configuration module monitors resource denials received by the virtual servers and determines if a virtual server is overloaded based upon the resource denials.” Ex. F at 2:52-55. Figure 1 depicts the dynamic resource configuration module 100 as including virtual server resource monitor 110, virtual server resource modifier 120, physical host load balancer 130 and dynamic virtual server mover 140. Figure 3, which also depicts the virtual server resource monitor 110, further includes virtual server 162B, recognizing that the virtual servers and physical hosts also necessarily make up part of the structure performing the claimed function.

Additionally, the specific description of VMware’s proposed “algorithm” is too narrow and adopting it would impermissibly narrow the term. For example, one of ordinary skill in the art would know from the claims and specification disclosure that there are many different ways to intercept systems calls, and resource denials can be accounted for and an overload determined in ways other than storing the denials in individual tables and using a pre-specified time window. *See* Akl Decl. at ¶ 37. The patentee recognizes this in several places throughout the specification. *See, e.g.*, Ex. F at 8:46-49; 8:65-66; 9:43-46. So, while IV does not dispute that virtual server resource modifier 110 is a part of the corresponding structure, the Court should also include the other components noted above, particularly dynamic resource modification module 100, and disregard VMware’s narrow recitation of the “algorithm.”

ii. “virtual server resource modifier” term, clause 5 of Ex. A

VMware again points to an overly narrow description of the structure performing the claimed function. The same arguments from Section (i) above apply to this disputed term. The virtual server resource modifier is certainly part of the disclosed structure, but the Court should also include dynamic resource configuration module 100 and the virtual servers in order to capture the full scope of the corresponding structure. One of ordinary skill in the art would understand that dynamic resource configuration module 100 is not a “black box” as VMware

contends, but rather, is a well-known component that plays a part in performing the claimed function.<sup>11</sup> *See* Akl Decl. at ¶ 38.

iii. “load balancing module” term, clause 6 of Ex. A

As with the other terms discussed above in sub-sections (i)-(ii), VMware’s proposal for this disputed term includes part of the structure that performs the claimed function but leaves out others. It ignores the fact that physical hosts 160A-C and virtual servers 162A-C are described as contributing to the claimed function (as seen in Figure 5 and the related description discussed above in Section 3(A)), as well as the fact that the physical host load balancing module is part of the dynamic resource configuration module 100, and includes the physical host resource monitor 540 and load balancing calculator 530, all of which should properly be considered as part of the corresponding structure. *See* Akl Decl. at ¶ 39.

iv. “dynamic virtual server mover” terms, clauses 7 & 8 of Ex. A

The corresponding structure for this disputed term is similarly broader than VMware’s description and many of the same arguments from Section 3(A) and sub-sections (i)-(iii) above are applicable here as well. Primarily among those arguments, VMware ignores the physical servers, virtual machines and dynamic resource configuration module 100, each of which are structural elements that play a part in performing the claimed function. *See supra* Section 3(A). For instance, the virtual server mover is part of the dynamic resource configuration module which is coupled to the physical hosts running the virtual machines. *See* Ex. F at Fig. 1 & 5; *see also*, Ex. F at 2:52-67. In addition, the specification describes physical host load balancer 130 as using the dynamic virtual server mover 140 to transfer virtual servers between hosts, indicating that the load balancer (also part of dynamic resource configuration module 100) is used, in part, to perform the claimed function. *Id.* at 5:9-13. Finally, one of skill in the art would understand from the specification’s disclosures that the ‘make, then break’ and ‘break, then make’ methods

---

<sup>11</sup> It’s also worth noting that as with all of its proposed constructions, VMware seems to incorporate its constructions of other terms, in this case, “virtual server overload signal,” “modify a resource allocation,” and “virtual server,” into the construction of the disputed term. This makes its proposed construction even more confusing, but most importantly, shows that it believes the structure of the disputed term at least includes the above identified terms, which belies its arguments in favor its narrowly defined structure.

of transfer are not the only methods for transferring virtual servers and thus the structure should not be so limited. *See* Akl Decl. at ¶ 40.

### **The '752 Patent**

United States Patent No. 7,949,752 (“the ’752 patent”) relates to the use and customization of network systems that allows delivery of highly reliable applications and services from the cloud to remote users. Through the use of novel virtualization systems and methods a client machine need only have basic browser support in order to have full use of cloud-based programs and applications. The network system and methods make this possible through virtualizing physical hardware into logical abstractions allowing resources to be distributed and managed at locations remote from end-user clients.

#### **4. Disputed Terms in the ’752 Patent**

##### **A. “exhausted” (’752 patent claims 1, 9 and 24)**

IV’s Proposed Construction	VMware’s Proposed Construction
“used up to the allotted or pre-determined amount”	“unavailable for reuse”

VMware claims that during the prosecution of the ’752 patent the patentee made a clear disclaimer regarding the scope of the disputed term, and therefore, the term should be construed according to its proposal of ‘unavailable for reuse.’ VMware then argues that in construing the term as IV proposes in prior litigation Judge Mitchell committed an error and misapplied the *Ecolab* case. Neither assertion is correct.

As Judge Mitchell found in the HCC Case, the prosecution history, read as a whole, does not contain a clear and unambiguous disclaimer with respect to the disputed term. While in response to a Section 102 rejection the patentee did state that the prior art failed to disclose a resource that is exhausted upon being consumed because it “teaches the exact opposite—i.e. that the ‘cartridges’ in Chou are reused,” the examiner disagreed, calling the argument “misleading” and further noting that “in the applicants invention, resources are reused.” *See* 2d DeBlois Decl., Ex. S at 13. The examiner then pointed the applicant to the exact portion of his specification that discloses reuse. *See* 2d DeBlois Decl., Ex. S at 19 (“On page 25 of the applicant’s specification

resources are defined as including processing time, memory storage space, and the like. While these resources may be at full capacity at any given time, the[y] can always be reused later they are not at full capacity.”). VMware misinterprets this as a mere ‘disagreement’ between the patentee and the examiner, noting that in the responsive office action the examiner withdrew the Chou reference and put forth another piece of prior art, “Cohn,” that formed the basis of his “new” Section 102 rejection. VMware then claims that Cohn disclosed “consuming non-reusable fees for using a service,” which proves that there was a disclaimer and the examiner agreed, at least passively. This is not entirely accurate, however, because VMware fails to mention two key facts; (1) the Chou reference was not withdrawn because the examiner somehow acquiesced to the patentee’s statement, and (2) the Cohn reference discloses reusable service resources as well.

With respect to the first issue, a careful analysis of the prosecution history reveals the flaw in VMware’s argument. In the same OA response where the patentee made the alleged disclaimer the applicant also amended then claim 86 to include the term “an amount of” the service resource being exhausted upon being consumed. *See* 2d DeBlois Decl., Ex. S at 2. Because Chou disclosed software containers called “cartridges” that the examiner equated to the claimed service resource, an ‘amount of’ the cartridge could never be exhausted regardless of the scope of the term ‘exhausted.’ Thus, the examiner’s withdrawal of Chou as a Section 102 reference was unrelated to the meaning of the disputed term.<sup>12</sup>

Regarding the second issue, a careful reading of the Cohn reference and file history shows that the reference also disclosed service resources that could be reused. In Cohn the examiner considered the service resource the “information providers” that comprise databases of information users can access and then be charged for based on level of access and the like. This is clearly a re-usable resource. Consider a scenario where subscriber A gets access to the database for information Y, information Y is not then “unavailable for reuse,” subscriber B can

---

<sup>12</sup> This makes logical sense considering that had the examiner truly believed that the patentee’s invention claimed the re-use of service resources—which he clearly did—as a matter of law he could not have withdrawn the objection and selected a reference that disclosed only non-reusable service resources simply to avoid the issue.

also subscribe to and receive access to information Y. VMware's claim that the Cohen reference disclosed "consuming non-reusable fees for using a service" thus evidencing a disclaimer in scope of the disputed term, is therefore incorrect.<sup>13</sup>

When the prosecution history and the references discussed therein are properly analyzed, it becomes apparent that, just as Judge Mitchell concluded, there was no clear disclaimer with respect to the disputed term. In fact, the evidence shows that the patentee recognized that the examiner was correct in his assertion that "exhausted" is broader than "unavailable for reuse" by abandoning that argument despite the fact that the Cohn reference disclosed reusable service resources. Instead, he amended the independent claims as suggested by the examiner to include the "URL" elements found in dependent claims. *Compare* 2d DeBlois Decl., Ex. S at 19-23 with Ex. S at 24-35. Therefore, VMware has not (and cannot) meet its burden of showing that there was a clear and unmistakable disclaimer to the scope of the disputed term.

VMware's dismissal of Judge Mitchell's logic in coming to the same conclusion regarding whether a disclaimer occurred is equally as flawed as its other arguments noted above. For example, VMware claims that Judge Mitchell committed an error in applying the *EcoLab* case. Judge Mitchell, however, considered this very same argument when it was made by HCC and specifically took note of it in her order, ultimately concluding that "[w]hile the facts in *EcoLab* may be distinguishable from the current case, its holding is nevertheless instructive. The prosecution history, when considered as a whole, demonstrates that the patentee did not make a clear and unmistakable disclaimer." *See* Ex. N at 33. Accordingly, the intrinsic record, a Primary Examiner at the USPTO, and a United States Magistrate Judge rejected VMware's position as should this Court.

---

<sup>13</sup> This characterization is also illogical because even if the reference was disclosing something similar to "consuming long-distance minutes" which VMware analogizes to, the service is the ability to make a long distance phone call and the service resource is the phone lines, switches, etc. that carry the phone call. Thus, even in VMware's characterization the service resource could never be "unavailable for reuse" since the phone lines and switches used to make long distance phone calls (regardless of the pricing structure for customers) are reused countless times.

**B. “consumed” (’752 patent claims 1, 9 and 24)**

IV’s Proposed Construction	VMware’s Proposed Construction
“used”	“used up”

VMware’s arguments as to why the disputed term should be limited to “used up” have no factual or legal basis. As an initial matter, VMware’s arguments here are a transparent attempt to save its carefully crafted non-infringement position regarding the ‘exhausted’ term previously discussed should the Court not find in its favor there. In addition, the idea that the disputed term was somehow “narrowed” during prosecution is equally unfounded. The claim amendment that VMware points to as support was made to specify that the agent is doing the consuming, not that “use” and “consumed” have two separate and distinct meanings. In fact, the portion of the file history that VMware cites to as support shows that the amendment was unrelated to the definition of the “use” or “consumed” terms, but rather, was made to specify that the agent is not a GUI as the examiner contended because a GUI cannot use the hardware resources of a server device. *See* VMware Br. Ex. 15 at 10; *see also* 2d DeBlois Decl., Ex. S at 38-40.

Furthermore, VMware’s argument that somehow the fact that ‘consumed’ is used in the past tense (as opposed to consume) indicates that it means “used up” overlooks the obvious point that the claim states “to be consumed,” which is the future tense. Thus, nothing has been used up, rather, something will be used. Similarly, VMware’s argument that there is a “definitional statement” in the specification defining ‘consumed’ as ‘used up’ also fails. Such a claim is belied by the fact that while the disputed term is referenced in a single instance as ‘used up’ it is referred to interchangeably with ‘used’ at least half a dozen times. *See* IV Br. at 5.

Finally, even the examples cited by VMware regarding voice mail and call placement are misplaced. VMware is correct if it is implying that the time a human spends making a phone call or leaving a voicemail cannot be reused, i.e., one cannot literally travel back in time, however, this is not what the patent is referring to. When the patent speaks of discrete units of a service resource being consumed while using a voicemail or long-distance telephone service the point is just that in one embodiment service resources can be structured in such a way that allows a



service provider to charge for each unit used. Either way, in no instance is the term limited to “used up,” a fact recognized by Judge Mitchell when facing the same arguments VMware makes now. *See* Ex. N at 29-31. Therefore, the Court should disregard VMware’s attempts to narrow the disputed term and adopt IV’s construction which is overwhelmingly supported by the intrinsic record.

**C. “service” (’752 patent claims 1, 3, 9 and 24)**

IV’s Proposed Construction	VMware’s Proposed Construction
“network functionality available to agents”	“an application that is used on behalf of a principal”

The Court should adopt IV’s proposed construction for this term because it is taken directly from the extensive specification disclosures, whereas VMware points to a single instance in which the disputed term is referenced as it proposes. There is no dispute that the patent discloses that a “service” could be compiled into a discrete program which the claimed agent could use based on an end-user’s (i.e. principal’s) instruction. *See* IV Br. at 7. But this is merely a single embodiment. Limiting the disputed term to one narrow embodiment mentioned a single time in the 35 pages of specification disclosure is inconsistent with basic claim construction principals. *See Comark Commc’ns, Inc.*, 156 F.3d at 1186-87. In this case, the overwhelming weight of the remaining disclosures indicate that the disputed term should be construed as IV proposes. For example, the specification notes that “an agent is operable to use a service.” *See* Ex. D at abstract. Examples of services include an email service, a voice mail service, etc., each supported by a “sub-system” of the network, i.e. network functionality, not applications. *Id.* at 13:21-28, Fig. 2. Further, the result of using a service might benefit a human user, but they are being used by the claimed agent. Therefore, the Court should adopt IV’s proposal because it captures the full scope of the disputed term as used in the intrinsic record and the claims.

**D. Means-Plus-Function Terms**

The disputed terms, as well as the Parties’ proposed function and structure can be found in Ex. B. The main dispute surrounds VMware’s attempts to limit its proposed structure by qualifying its identification with the phrase “as described in,” followed by citations to a narrow

embodiment of such structure. For the reasons set forth in IV's Opening Brief, Section 3(A)(4), this should not be allowed.

VMware argues that § 112 ¶ 6 requires that it include specific specification citations that correspond to the disclosed structure, but this is not so. This Court has previously construed means plus function terms where the corresponding structure for performing the claimed function includes only the disclosed structure, e.g., 'widget 26,' and § 112 ¶ 6 requires nothing more. *See, e.g., VLSI Tech. LLC, v. Intel Corp.*, Case No. 1:19-CV-00977-ADA, Claim Construction Order, Dkt. No. 101. In any case, the real issue here is that VMware's citations exclude disclosed embodiments specifically linked to performing the claimed function and are unnecessary. Accordingly, the Court should adopt the proposed structure that the Parties agree on and not read in VMware's narrowing citations.

The only other dispute regarding the '752 patent's means plus function terms is with respect to the structure of "means for monitoring an amount of the service resource used by the network-based agent." The Parties' proposals are below:

IV's Proposed Construction	VMware's Proposed Construction
<u>Function:</u> monitoring an amount of the service resource used by the network-based agent" <u>Structure:</u> Service Wrapper 26	<u>Function:</u> monitoring an amount of the service resource used by the network-based agent" <u>Structure:</u> monitor as described in '752 patent at 16:50-61.

The Court should adopt IV's proposed structure because it captures the full scope of the specifications disclosure whereas VMware's proposal includes only a sub-component of the service wrapper 26 which performs the function as a whole. *See* IV Br. at 8. Furthermore, VMware's argument that the converter sub-component of service wrapper 26 is irrelevant to performing the claimed function is misplaced. Since the converter converts the computer language used by a service to the computer language used by the agent server, which in turn interacts with monitor 50 to enforce service resource usage of the agents when using the service, the converter does perform part of the claimed function. In other words, without the converter it would not be possible to monitor and enforce service resource usage since the components using the services and service resources would be unable to communicate with the components

monitoring the same. *See* Ex. D at 16:22-61. Therefore, the Court should adopt IV’s construction as it accurately reflects the scope of the corresponding structure.

### **The ’051 Patent**

The ’051 patent teaches combinations of virtualized components, tunnels, and physical interfaces that collectively allow enterprise or data center operations to be siloed on a per-tenant basis so that each tenant may use the same physical hardware as another tenant without risk of security breach or hardware failure. And, because the tenant’s servers are virtualized, the addresses of one tenant’s virtual servers may overlap with the addresses of another tenant’s virtual servers. Such virtual isolation thus reduces the risk of a security breach, while continuing to provide tenants with a lot IP addresses.

#### **5. Disputed Terms in the ’051 Patent**

##### **A. “virtual server” (claims 1, 3, and 6)**

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning Alternatively – “virtual machine(s) that reside(s) on a physical server and use(s) the physical server’s resources, but [has/have] the appearance of being a separate, dedicated machine(s)”	“a process executing on a host computer that accepts communications requests”

Notably, in the four pages dedicated to this term in its Opening Brief, VMware does not discuss IV’s primary proposal, that “virtual server” should be construed with its plain and ordinary meaning. As explained in IV’s Opening Brief, plain and ordinary meaning is the correct construction because the ’051 patent uses this well-understood term in a non-limiting manner. IV Br. at 41-42. At the time of the inventions of the ’051 patent, a POSITA would readily understand that the term meant a virtualized server and no further definition is necessary. *See* Declaration of Dr. Vijay Madiseti (hereinafter “Madiseti Decl.”) at ¶¶ 31-33. As such, “virtual server” as used in the patent, has an even broader meaning than even IV’s alternative construction. IV’s alternative proposed construction, the entire focus of VMware’s Opening Brief, is meant to define the term in the context of known virtualization practices. In no event should the term be defined as VMware proposes. VMware’s construction finds no support in the

intrinsic record and runs counter to how a POSITA would have understood the term at the time of the invention.

Finding no support in the '051 patent's written description, VMware resorts to citations from a patent application incorporated by reference into the '051 patent's specification, U.S. Patent App. Ser. No. 09/432,286 ("the '286 application"). Not only does the '286 application's specification not define, or otherwise disclaim the scope of, virtual server in the manner advocated by VMware, but the citations relied on by VMware are describing a physical host server, not a virtual server. *See* VMware Br. at 34-38 (citing 6,976,258) (distinguishing "a server, which is a process, executing on a dedicated physical host" from "a virtual host server [that] services numerous client requests for multiple virtual hosts" in same paragraph); Abstract (differentiating "plurality of virtual hosts . . ." from ". . . on a single physical host computer"); 1:16-18 (referencing "a server program executing on a single physical host computer" indicating that servers, virtual or otherwise, are not exclusively referred to using "process" language). In addition to there being no teaching or suggestion in the '286 application that the "process" characterization of a server was intended to be limiting, there is no disclosure in the '051 patent purporting to limit the meaning of virtual server to particular examples or embodiments discussed in the '286 application. *See* Ex. H at 3:64-67 (characterizing incorporated reference as example disclosing "a method for creating such a private virtual server", but not expressly defining a virtual server as such).

Further, VMware's proposed construction fails to account for the "virtual" aspect of the disputed "virtual server" term. As discussed in IV's opening brief, VMware's proposed construction of "a process executing on a host computer that accepts communications requests" does not contain any wording that would enable a POSITA to distinguish the claimed virtual server from a networking-related process associated with a conventional, physical server.<sup>14</sup> It is well-established that claim construction "begins and ends in all cases with the actual words of

---

<sup>14</sup> Indeed, this is precisely because the "process" language from '286 application that VMware has latched onto was language describing an example of a physical, non-virtual server.

the claim.” *Renishaw PLC v. Marposs Societa' per Azioni*, 158 F.3d 1243, 1248 (Fed. Cir. 1998). VMware’s proposed construction violates this fundamental principle by disregarding the term “virtual.”

Although IV believes a construction beyond plain and ordinary meaning is unnecessary, the proposed alternate construction seeks to clarify the “virtual” aspect of a “virtual server” by way of an embodiment of a virtual server pertinent to this case. The specification makes it clear that a virtual server “can service requests to multiple network addresses or domain names [t]hus, the functionality of numerous hosts is provided by a single physical host computer” and further describes the inability of existing virtual servers “to allocate appropriate amounts of computer resources of the physical host computer to servicing client requests made to specific virtual hosts.” Ex. H at 2:50-64, 3:57-64. In light of such disclosure in the specification and its known meaning in the art, a POSITA would readily have understood a virtual server as an example of virtual machine that abstracts physical server resources and appears from the outside to be a fully-functional, dedicated physical server. *Id.* at 2:60-62 (“ . . . the virtual host server can service requests to multiple network addresses or domain names. Thus, the functionality of numerous hosts is provided by a single physical host computer.”). *See Medrad, Inc. v. MRI Devices Corp.*, 401 F.3d 1313, 1319 (Fed. Cir. 2005) (it is “entirely proper to consider the functions of an invention in seeking to determine the meaning of particular claim language”); *see also* Madisetti Decl. at ¶¶ 31-33.

**B. “physical interface[s]” (claims 1 and 3)**

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	“hardware that provides a point of communication between two or more devices”

The disputed term “physical interface” should be construed as its plain and ordinary meaning because it is a well-understood construct to those of skill in the art and requires no construction other than plain and ordinary meaning. The only exceptions to giving a term its plain and ordinary meaning are when the patentee has acted as his own lexicographer or made a clear and unambiguous disavowal of scope. *See Thorner*, 669 F.3d at 1365. Neither of those

exceptions apply here. For example, the disputed term is disclosed in the specification as one of ordinary skill in the art would expect and understand to be the plain and ordinary meaning of the term. *See* Ex. H at 10:37-42; *see also* Madisetti Decl. at ¶¶ 34-36.

VMware's construction improperly reads in limitations inconsistent with the description of the disputed term in the specification. Specifically, VMware's limitation "between two or more devices" is contrary to specification's broader disclosure. *See* Ex. H at 10:37-38.

Accordingly, IV's proposed construction of plain and ordinary meaning should be adopted.

**C. "storing . . ." / "storing . . ." / "receiving . . ." / "determining . . ." / "determining . . ." / "sending . . ." / "using . . ." (claims 1 and 3)**

VMware characterizes these claim elements as "ambiguous" and having a "significant antecedent basis problem that makes the claims impossible to decipher" in view of the constituent terms "physical interface(s)" and "tunnel identifier(s)" recited in claims 1 and 3. VMware offers no further clarification to support its assertion nor any statements explaining how or why specific instances of the constituent terms lack antecedent basis either individually or in combination.

There are no antecedent basis issues in claims 1 and 3. The disputed claims each begin by reciting generalized (i.e., unqualified by an article such as "a", "an", or "the") plural forms of each constituent term: "physical interfaces" and "tunnel identifiers." Ex. H at claims 1 and 3. Next, the claims sequentially recite the constituent terms using an indefinite article ("a") followed by a definite article ("the"): "a physical interface" / "the physical interface" and "a tunnel identifier" / "the tunnel identifier." *Id.* Lastly, the claims sequentially recite the constituent terms using an indefinite article followed by a qualified definite article that refers directly back to the preceding indefinite article: "determining . . . a physical interface and tunnel identifier" / "the determined [physical interface/tunnel identifier]" (claim 1), and "using . . . to identify a physical interface and tunnel identifier" / "the identified [physical interface/tunnel identifier]." *Id.*

As a matter of grammar, articles identify particular known (definite form) or unknown (indefinite form) instances of the terms they qualify. It is well-established under antecedent basis law that claim terms qualified by a definite article are interpreted to refer back to preceding instances of the term qualified by an indefinite article. *Baldwin Graphic Sys., Inc. v. Siebert, Inc.*, 512 F.3d 1338, 1342-43 (Fed. Cir. 2008) (recognizing “an indefinite article ‘a’ or ‘an’ in patent parlance” and “the subsequent use of definite articles ‘the’ or ‘said’ in a claim to refer back to the same claim term”). Further, reciting a term qualified by an indefinite article after a previous recitation of the term in generalized plural form does not necessarily implicate the singular term as an instance of the plural term, particularly when additional qualifying language is present. *See Medtronic, Inc. v. Guidant Corp.*, 465 F.3d 1360, 1377 (Fed. Cir. 2006) (finding no requirement that “a control signal” necessarily count as instance of previously recited “cardiac signals” and noting distinct qualifying language “control” and “cardiac”). The literal claim language as written is therefore entirely consistent with these established principles of antecedent basis law.

In addition to antecedent basis being satisfied on its face, the surrounding context of the claim language expressly clarifies which instances of “physical interfaces” and “tunnel identifiers” are associated with “customer lookup table/information” and the which instances are associated with “customer forwarding table(s)/information.” The claims when read in their entirety, in conjunction with antecedent basis being satisfied on its face, suffice to describe to a POSITA what the inventors were claiming. *See ACTV, Inc. v. Walt Disney Co.*, 346 F.3d 1082, 1088 (Fed. Cir. 2003).

As explained in IV’s Opening Brief, plain and ordinary meaning is the correct construction because the ’051 patent uses well-understood claim terms combined in a clear and straightforward manner. *IV Br.* at 45-46. At the time of the inventions of the ’051 patent, a POSITA would readily have understood the intended meaning of each disputed term. Thus, no further definition is necessary. *See Madisetti Decl.* at ¶¶ 37-39.

Despite there being no need to construe these omnibus claim elements beyond their plain and ordinary meaning, VMware characterizes its proposed constructions as “a simple way to resolve this problem that is fully supported by the specification.” This is not true and precisely why VMware’s proposed constructions are inappropriate i.e., because they violate the well-settled rule of claim construction that limitations described in particular embodiments in the specification should not be imported into a claim unless accompanied by statements of clear disavowal or otherwise required by the specification. *See Arlington Indus., Inc. v. Bridgeport Fittings, Inc.*, 632 F.3d 1246, 1254 (Fed. Cir. 2011).

VMware cites only examples in the specification where an incoming and an outgoing directionality is discussed. *VMware Br.* at 41-42. However, the specification confirms that the claim is broader than these examples. Indeed, the specification contemplates numerous instances where the claimed tunnel switching activity can operate in either or both directions. *See e.g.*, Ex. H at 12:6-20 (“in the opposite direction”); 11:51-54 (“tables 800 and 900 operate to switch transmissions in both directions”); 14:20-21 (“the tunnel functions as a bi-directional data pipe”); Figs. 6-7; 10B-11 (using bidirectional arrows and/or arrows in both directions to depict tunnel switching activity). Thus, VMware improperly asks the Court to read preferred embodiments out of the claim. *See Accent Packaging, Inc. v. Leggett & Platt, Inc.*, 707 F.3d 1318, 1326 (Fed. Cir. 2013) (“a claim interpretation that excludes a preferred embodiment from the scope of the claim is rarely, if ever, correct.”).

Other than its conclusory antecedent basis allegation concerning “physical interface(s)” and “tunnel identifier(s)” in particular, VMware offers no explanation as to why a POSITA would fail to understand any of the constituent claim terms recited in the elements it proposes to construe, nor why a POSITA would fail to understand how the combinations of said claim terms interrelate with each other. VMware has not met its burden of showing that the plain and ordinary meaning of the literal claim language is ambiguous or contradictory to the specification in any way.



**D. “customer forwarding [table(s)/information]” (claims 1 and 3)**

IV’s Proposed Construction	VMware’s Proposed Construction
“table(s) containing [a set/sets] of customer specific forwarding information” / “set(s) of customer specific forwarding information”	See construction of "storing a customer lookup table, the customer lookup table storing associations between physical interfaces and tunnel identifiers identifying tunnels for private networks and a plurality of customer forwarding tables" / “storing . . . customer forwarding information, . . . the customer forwarding information associating network addresses with physical interfaces and tunnel identifiers"

VMware does not specifically address IV’s proposed construction of the terms “customer forwarding table(s)” and “customer forwarding information” recited in claims 1 and 3.

Throughout much of the preliminary claim construction and subsequent meet-and-confer process, VMware’s original position was that it was relying on its separately proposed constructions for the terms discussed *supra* in Section 5(C) to construe the disputed terms at issue here (those separate, larger claim terms recite instances of the disputed terms within). However, VMware’s revised constructions no longer propose to alter or qualify the meaning of “customer forwarding table(s)” and “customer forwarding information” in any way. Rather, they focus on alleged antecedent basis issues associated with the “physical interface” and “tunnel identifier” terms. Although the parties did not specifically discuss dropping or agreeing upon a construction for this term, it appears VMware is suggesting by omission that it does not object to IV’s proposed constructions.

IV’s proposed constructions closely track the underlying claim language and seek only to add the word “set” or “sets” where appropriate to help one differentiate between singular versus plural instances of the disputed claim terms, and further to maintain consistency between claims 1 and 3. In addition to closely mirroring the literal claim language, IV’s constructions are consistent with a POSITA’s understanding of the terms in view of the specification. For example, the ’051 specification explains:

FIG. 8 is an embodiment of a customer lookup table **800**, and FIG. 9 is an embodiment of a set of customer forwarding tables **900**. Together, table **800** and set of tables **900** are suitable for switching a set of transmissions . . . Customer lookup table **800** is used as an

index into the correct customer forwarding table from the set of customer forwarding tables 900.

Ex. H at 11:42-54; Figs. 8-9; *see also* Madisetti Decl. at ¶¶ 40-41.

### **The '818 Patent**

The inventions of United States Patent No. RE 44,818 (“the '818 patent”) provide improved server virtualization solutions in demanding computing environments, such as data centers. Specifically, the '818 patent teaches the dynamic managing of networks, virtual machines, virtual storage, and virtual input/output resources using a more granular and flexible quality-of-service scheme that includes a hierarchical classification and enforcement mechanism for allocating bandwidth within those networks.

#### **6. Disputed Terms in the '818 Patent**

##### **A. “hierarchical token bucket resource allocation” / “token(s)” (claims 1, 17, 30, 32 and 42)**

IV's Proposed Construction	VMware's Proposed Construction
Plain and ordinary meaning	The specific class-based scheduling algorithm known in the art as the “hierarchical token bucket” / “token” as used in “hierarchical token bucket resource allocation”

Here, VMware seeks to cabin the well-understood concept of a hierarchical token bucket algorithm into a specific, extrinsic definition contrary to the intrinsic record. It is improper for VMware to substitute extrinsic evidence inconsistent with the broad disclosure found in the intrinsic record. *See Bell & Howell Document Mgmt. Prods. Co. v. Altek Sys.*, 132 F.3d 701, 706 (Fed. Cir. 1997) (“reliance on extrinsic evidence to interpret claims is proper only when the claim language remains genuinely ambiguous after consideration of the intrinsic evidence”).

As discussed in IV's Opening Brief, the specification supports reading the term in accordance with its plain and ordinary meaning as understood in the context of the hierarchical QoS processes described therein. *See e.g.*, Ex. G at 2:1-18; 8:20-44; 9:61-10:67; 11:25-27; 11:51-52; 13:21-23; 13:40-44; Figs. 12-13. And, contrary to VMware's proposed construction, the intrinsic record does not limit the term, let alone define it with a particular, extrinsic definition. Accordingly, IV's proposed construction of plain and ordinary meaning is proper

since a POSITA would readily understand the intended meaning of the disputed term in view of the specification. *See* Madisetti Decl. ¶¶ 42-45.

As to VMware's limiting extrinsic construction, the inventors do not specifically cite or reference any extrinsic definition of a hierarchical token bucket scheme in the intrinsic record. Further, the claim drafters chose to expressly recite the term as using the indefinite article "a" ("a hierarchical token bucket resource allocation") as opposed to using definite phrasing such as "Hierarchical Token Bucket resource allocation" indicating that there was no intent to limit the definition to a specific, extrinsic approach or term of art. VMware asserts that "HTB is a "proper noun" that refers "specifically to the scheduling mechanism known as 'hierarchical token bucket' and not anything else" and points to the following language from the '818 patent: "using scheduling and queuing methods such as hierarchal token bucket (HTB)." Ex. G at 9:64-65. However, VMware's argument appears to collapse the distinction between proper nouns and acronyms. Although acronyms are always capitalized, many acronyms do not stand for proper nouns and should not be capitalized in their written-out form. As is the case in the '818 patent, the specification first and foremost identifies "hierarchical token bucket" without capitalization as a common noun. The fact that the "(HTB)" acronym which follows uses capitalization is irrelevant because it is a general property of all acronyms.

Further, VMware's arguments contradict its own position. As it acknowledges in its Opening Brief, "HTB was invented or at least popularized by Martin Devera when he implemented it as a queuing discipline in the Linux operating system." VMware Br. at 42. VMware thus concludes that the term could only logically refer to "the very specific method of allocating bandwidth resources referred to in the art." *Id.* However, as evidenced by VMware's own extrinsic evidence citations, numerous implementations of a hierarchical token bucket algorithm have been introduced into the art since the initial usage of the concept as a Linux queuing mechanism. *Id.* (referencing Snoeren Decl. ¶¶ 107-112) (listing multiple hierarchical token bucket implementations including Linux implementation, WLAN implementation, container implementation, etc.). It is well-settled that the meaning of a disputed term (as

understood by a POSITA) is measured relative to the critical date of the invention at issue, not the inception date of the term in the prior art. *See Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.*, 381 F.3d 1111, 1116-17 (Fed. Cir. 2004) (“[a] court construing a patent claim seeks to accord a claim the meaning it would have to a person of ordinary skill in the art at the time of the invention.”).

Although the concept of a hierarchical token bucket algorithm may have originated and first been implemented as a specific Linux-based queuing mechanism, various additional implementations emerged over the years prior to the invention of the ’818 patent demonstrating that usage of the term in the art evolved and diversified over time following its initial inception. Even assuming, *arguendo*, that VMware’s cited reference was the way to implement a hierarchical token bucket when it was written, the patent’s broad disclosure demonstrates that it was not the only way at the time of invention. Accordingly, VMware’s characterization of the term as a “very specific method . . . referred to in the art” is misguided and aimed solely at narrowing the meaning of the term without specifically identifying any deficiencies in the claim language itself or associated passages in the specification. The appropriate construction of “hierarchical token bucket resource allocation” should be its plain and ordinary meaning as interpreted in light of the specification.

**B. “enforcing . . .” / “receiv[e/ing] . . .” / “classify[ing] . . .” / “compar[e/ing] . . .” / “forward[ing] . . .” / “buffer[ing] . . .” (claims 1, 17, 30, 32, 33, 37-39, 42)**

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	“enforcing . . . across the physical [storage network] interface of the virtual I/O server” / “receiv[e/ing] in the virtual I/O server” / “classify[ing] in the virtual I/O server” / “compar[e/ing] in the virtual I/O server” / “forward[ing] in the virtual I/O server” / “buffer[ing] in the virtual I/O server”

Here, VMware asks the Court read in the limitation “virtual I/O server” into nearly every element of claims 1, 17, 30, 32, 33, 37-39, and 42 of the ’818 patent despite those words not appearing in the body of any asserted claim.

As support, VMware primarily relies on case law standing for the proposition that limitations that “repeatedly and consistently” describe the invention may limit claim scope. VMware Br. at 45. The facts here, however, are much different. In both *SkinMedica* and *Microsoft*, the patentee made statements describing two embodiments of a single term as necessarily distinct resulting in the court finding that the disputed term, in the context of the claim, excluded one of those embodiments.<sup>15</sup>

VMware’s proposed construction, however, does not seek to merely narrow the functional scope of an existing claim term based on its limiting characterization in the specification. Rather, it proposes to import wholesale new structural limitations into nearly every element of every claim. This violates the well-established precedent that examples from the specification should not be imported into the claim where the claim language is not limiting and/or is broader than embodiment(s) described in the specification. *See Arlington Indus., Inc.*, at 1254 (Fed. Cir. 2011).

This is particularly true when the importation of structural limitations is at issue, as is the case here. *See Paragon Sols., LLC v. Timex Corp.*, 566 F.3d 1075, 1084-85 (Fed. Cir. 2009) (rejecting importation of structural limitation into claims requiring “electronic positioning device” and “physiological monitor” to include separate physical structures; district court erroneously relied on examples in spec that separated structures as limiting).

Further, with respect to method claims 1, 32, and 42, they each recite “A method of facilitating management of input/output subsystems in a virtual input/output server, comprising” in the preamble. It is well-settled, however, that limitations from the preamble should not be imported into the body of the claim if the claim body defines a structurally complete invention. *See Georgetown Rail Equip. Co. v. Holland L.P.*, 867 F.3d 1229, 1236 (Fed. Cir. 2017) (preamble language not limiting if claim and spec otherwise recite structurally complete invention).

---

<sup>15</sup> *C.f. SkinMedica, Inv. v. Histogen, Inc.*, 727 F.3d 1187, 1203-04 (Fed. Cir. 2013); *Microsoft Corp. v. Multi-TechSystems, Inc.*, 357 F.3d 1340, 1347-48 (Fed. Cir. 2004).

Furthermore, it is entirely proper to claim method body elements as a series of functional steps. *See Cox Commc 'ns, Inc. v. Sprint Commc 'n Co. LP*, 838 F.3d 1224, 1232 (Fed. Cir. 2016) ("the asserted claims are method claims, so it makes sense to define the inventive method as a series of functions").

VMware does not allege that any of the leading terms “enforce[e/ing],” “receiv[e/ing],” “classify[ing],” “compar[e/ing],” “forward[ing],” and “buffer[ing]” are in dispute, nor does it specifically identify any other literal claim terms as allegedly unclear or expressly redefined by the specification. This is precisely because the plain and ordinary meaning of the disputed terms is clear. All VMware seeks to do is narrow the claims by importing unclaimed structure from embodiments in the specification. This does nothing to clarify the meaning of the literal claim terms at issue other than imposing extraneous, unrecited limitations in nearly every claim element. *See, e.g., O2 Micro*, 521 F.3d at 1362 (district courts not required to construe every limitation present in asserted claims and should focus on claim terms for which the parties have presented a fundamental dispute).

The terms at issue are used clearly and unambiguously throughout the claims in a manner that is consistent with the specification. Thus, a POSITA would properly understand these terms and accord them their plain and ordinary meaning. *See Madisetti Decl.* at ¶¶ 46-47.

**C. “maintain[ing] a connection over a network fabric” (claims 1, 17, 30, 32, 42)**

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	"maintaining a connection between the physical interface of the application server and the physical interface of the virtual I/O server over a network fabric"

VMware argues that two additional limitations should be read into the disputed term including “the physical interface of the application server” and “the physical interface of the virtual I/O server.”

In support of its desire to read in the limitation “and the physical interface of the virtual I/O server,” VMware first references the opinion of its expert Dr. Snoeren in stating that “a POSITA would understand that a physical interface connects to a network fabric” and further

concluding that a physical interface limitation must be read into the claim. It is well-settled that limitations from the specification should not be read into a claim unless the patentee acts as his or her own lexicographer or the specification otherwise makes clear statements of disavowal. *See Thorner*, 669 F.3d at 1365; *Arlington*, 632 F.3d at 1254. Dr. Snoeren’s own statement contradicts VMware’s argument, since he relies directly on language from the ’818 patent specification stating that “[i]n one implementation, the I/O fabric driver stack and I/O fabric PHY interface can employ the Reliable Connections (RC) supported by the Infiniband standard.” Ex. G at 3:9-11. As expressed literally by the patentees, the statement at issue merely describes one implementation and was clearly not intended to be limiting. Further, the statement appears to be describing characteristics of an exemplary I/O server as opposed to a network fabric itself.

Next, to support reading in the additional limitation “between the physical interface of the application server,” VMware argues that “Figure 2 depicts the protocol stack of an application server” and again references the opinion of its expert Dr. Snoeren in concluding that “the physical interface of the virtual I/O device connects to the network fabric.” VMware Br. at 47. Once again, VMware mischaracterizes the specification to suggest that Fig. 2 discloses “the protocol stack of an application server” where in fact the ’051 specification states “(“virtual network interface, in one implementation, emulates an Ethernet NIC [and] plugs in at the bottom of the network stack.” Ex. G at 4:9-13; *see also* at 4:2-4 (stating “[i]nserted into the network and storage protocol stacks are virtual interface drivers configured to intercept storage and network I/O messages” thus further contemplating the existence of multiple implementations of a protocol stack).

Also, despite the disputed term’s inclusion of “maintain[ing] a connection over a network fabric,” VMware’s argument focuses on the “network fabric” portion of the term in isolation. VMware further overlooks the presence and significance of the claim language that directly follows and qualifies the disputed term: “. . . to a virtual [storage/network] interface layer of an application server.” Taking the entire claim element in context, the patentees have already defined structural relationships between “a connection,” “a network fabric,” and “a virtual

interface layer of an application server.” Thus, it is unclear why it would be necessary or permissible to impose additional “physical interface” limitations when the claim drafters have already defined structure relating “a connection” to “a virtual interface.” It would have readily been understood by those having ordinary skill in the art in view of the specification that the term “connection” generally is not limited to a “physical connection” since connections can just as readily be logical or virtual, for example. Ex. G at 4:9-13 (“virtual network interface, in one implementation, emulates an Ethernet NIC [and] plugs in at the bottom of the network stack”); 6:29-32 (“[v]irtualization software in the virtual machine server abstracts the underlying hardware by creating an interface to virtual machines, which represent virtualized resources”); *see also* Madisetti Decl. at ¶¶ 48-49.

**D. “virtual [network/storage network] interface layer of an application server” (claims 1, 17, 32, 42)**

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	“a virtual storage network interface to higher layers of the virtual node in an application server”/ “virtual network interface layer to higher layers of the virtual node in an application server” / “virtual interface layer to higher layers of the virtual node in an application server”

VMware’s proposed construction introduces two additional limitations to these terms: (1) “. . . to higher layers;” and (2) “. . . of the virtual node.” In doing so, VMware asks the Court to ignore two claim construction canons by asking the Court to read in preferred embodiments and to ignore the patentee’s use of the limitation in some claims but not others.

The dispute concerning “. . . to higher layers” narrows to VMware’s proposal to impose a plural “to higher layers” requirement in an attempt to read in limitations from the specification. Ex. G at 4:29-32 (“FIG. 2 illustrates the protocol stack and modules of an application server 102 according to one possible implementation of the invention.”). This disregards the well-established principle that a claim term is not limited to specific or even preferred embodiments of the specification absent a clear disclaimer or disavowal of claim scope. Worse, as described in IV’s Opening Brief, VMware’s construction would read out preferred embodiments. Ex. G at



4:14-32; 4:51-62; Fig. 2; *see also* IV Br. at 30. Such a construction is “rarely, if ever, correct.” *See Accent Packaging*, 707 F.3d at 1326.

With respect to VMware’s second limitation, “virtual node,” that term is recited in certain claims but not others. For instance, it appears in claims 32 and 42 but is absent from claims 1 and 17. Thus, the applicant knew how to claim “virtual node” when desired and how not to claim it when the intent was to not use the term. Underscoring this point, in some instances the virtual node language even directly qualifies the disputed claim term, such as in claim 42 reciting “wherein the virtual storage network interface layer is associated with a virtual storage node identifier.” In light of such explicit usage elsewhere in the claims, the patentees established a clear familiarity with the concept of a virtual node and further displayed the ability to recite said limitation expressly in the claims when intended. VMware’s proposed inclusion of the words “of the virtual node” as qualifying the virtual interface layer(s) is, therefore, another attempt to read additional limitations into a claim term contrary to the surrounding context. *See Hockerson–Halberstadt*, 183 F.3d at 1374 (“the context of the surrounding words of the claim also must be considered in determining the ordinary and customary meaning of those claims.”).

VMware further fails to even suggest that either constituent phrase: “virtual interface layer . . .” or “. . . of an application server” is itself ambiguous, or that the relationship between the two phrases is ambiguous, which further undermines their position. To the extent the ’818 specification does discuss the concept of nodes, there are examples in which application servers are referred to as instances of “nodes” generally, but not “virtual” nodes as VMware would suggest. *See e.g.*, Ex. G 5:37-39 (“For example, the monitor module 250 is operative to automatically discover nodes (e.g., other application servers 102, virtual I/O servers 106).” For at least this additional reason, VMware’s proposed inclusion of the language “of the virtual node” is inappropriate as it deviates from the literal words of the disputed claim language at issue and seeks only to read in limitations that are neither required nor suggested by the intrinsic record. Again, VMware does not even argue that the constituent phrases are vague either individually or in combination. *See Homeland Housewares*, 865 F.3d at 375 (“[t]he claim

construction inquiry . . . begins and ends in all cases with the actual words of the claim.”); *Renishaw*, 158 F.3d at 1248 (“[t]his is so because the claims define the scope of the right to exclude”).

**E. Alleged Means Plus Function Elements (claim 17) (individually set forth in Ex. C)**

VMware alleges that ten elements of claim 17 implicate § 112 ¶ 6 and are also invalid for failure to disclose sufficient structure. VMware Opening Br. at 49-50. In its brief, however, VMware dedicates less than a page to supporting these allegations, all of which are conclusory and unsupported by the intrinsic record. As discussed in detail in IV’s Opening Brief, Section C(5), these elements of claim 17 (listed in full in Ex. C) are not subject to § 112 ¶ 6 because: (1) “means for” is not present in any claims, (2) the prefatory language used does not consist of merely a nonce word followed by function, and (3) the claim language itself provides sufficient structure to avoid the application of § 112 ¶ 6. *See Fisher-Rosemount Sys.*, 2019 WL 6830806, at \*15-\*16.

VMware first claims that ‘module’ is a nonce word and the preceding qualifier ‘input/output virtualization’ term “fails to impart any structural significance.” This argument, however, overlooks the significance of the term ‘input/output virtualization module’ as a whole, the language of the claim, and the knowledge of one of skill in the art. For instance, an ‘input/output virtualization module’ connotes a distinct structure to one of skill in the art. *See Madisetti Decl.* at ¶ 50. The fact that in some cases when a module is claimed alone without any qualifiers it can be considered a nonce word is not applicable here because in claim 17 the term is qualified as an ‘input/output virtualization module.’ As the court in *Zeroclick* noted, it is inappropriate to apply § 112 ¶ 6 to an alleged nonce term when it is qualified with terms denoting structure to one of skill in the art. *See Zeroclick, LLC v. Apple Inc.*, 891 F.3d 1003, 1007-09 (Fed. Cir. 2018). While the claim at issue in *Zeroclick* was “user interface code,” the principal applies to the term ‘input/output virtualization module’ as well, because ‘input/output

virtualization’ is understood as a structural element to one of skill in the art, particularly in light of the claim read as a whole. *See* Madisetti Decl. at ¶ 50.

VMware’s argument that the claim elements lack structural support is further belied by the claim language preceding the ‘input/output virtualization module’ elements. The claim language itself includes “an input/output fabric interface,” “a storage network interface,” “one or more processors,” and “a memory.” *See* Ex. G at claim 17. Thus the ‘input/output virtualization module’ is tied—without leaving the four corners of the claim—to the above noted structure as understood by one of ordinary skill. *See* Madisetti Decl. at ¶ 50. This also indicates that application of § 112 ¶ 6 is not appropriate. *See Telcordia Techs., Inc. v. Cisco Sys., Inc.*, 612 F.3d 1365 (Fed. Cir. 2010).

Finally, as with the alleged means-plus-function terms of the ’686 patent, VMware here misapplies the standard for whether corresponding structure is disclosed. Assuming that the terms are found to be subject to § 112 ¶ 6—which they should not be—the specification does in fact disclose corresponding structure for these disputed terms, which IV noted in its Opening Brief at pg. 36-37, and its disclosure of proposed constructions (*see* Ex. C). For instance, the function claimed in element 1 of Ex. C is tied to structure in the specification including a network fabric, and a virtual storage network interface layer. The same holds true for elements 2-3 which are clearly linked to the structure of a physical network storage interface. *See* Ex. G at 3:4-7; Fig. 7. The remainder of the disputed terms (elements 4-10) either individually recite sufficient structure or refer back to the structural elements discussed above. *See, e.g.*, Ex. C (defining structure for disputed terms). This fact would be easily understood by one of skill in the art, which is the standard for the sufficiency of structure in a § 112 ¶ 6 analysis. *See Telcordia Techs.*, 612 F.3d at 1377 (“claim definiteness depends on the skill level of an ordinary artisan. Therefore, the specification need only disclose adequate defining structure to render the bounds of the claim understandable to an ordinary artisan.”); *see also, Intelligent Automation Design*, 2020 WL 486830, at \*3-\*5.

### **Conclusion**

For the reasons stated herein and in IV's Opening Claiming Construction Brief, IV's proposed constructions should be adopted.

Dated: March 27, 2020

/s/ Robert R. Gilman  
DEREK GILLILAND  
State Bar No. 24007239  
SOREY, GILLILAND & HULL, LLP  
109 W. Tyler Street  
Longview, Texas 75601  
903.212.2822 (telephone)  
903.212.2864 (facsimile)  
derek@soreylaw.com

KARL RUPP  
State Bar No. 24035243  
NIX PATTERSON L.L.P.  
1845 Woodall Rodgers Fwy. Suite 1050  
Dallas, Texas 75201  
972.831.1188 (telephone)  
972.444.0716 (facsimile)  
krupp@nixlaw.com

#### **OF COUNSEL:**

Paul J. Hayes  
phayes@princelobel.com  
Matthew Vella  
mvella@princelobel.com  
Robert R. Gilman  
rgilman@princelobel.com  
Jonathan DeBlois  
jdeblois@princelobel.com  
Alex Breger  
abreger@princelobel.com  
PRINCE LOBEL TYE LLP  
One International Place, Suite 3700  
Boston, MA 02110  
Tel: (617) 456-8000  
Fax: (617) 456-8100

#### **COUNSEL for PLAINTIFFS**

CERTIFICATE OF SERVICE

I hereby certify that on the 27<sup>th</sup> day of March, 2020, I electronically filed the foregoing with the Clerk of using the CM/ECF system which will send notification of such filing to all counsel of record.

/s/ Robert R. Gilman

**UNITED STATES DISTRICT COURT  
WESTERN DISTRICT OF TEXAS  
AUSTIN DIVISION**

Intellectual Ventures I LLC and  
Intellectual Ventures II LLC,

Plaintiffs,

v.

VMware, Inc.,

Defendant.

Civil Action No. 1:19-cv-01075-ADA

**PLAINTIFFS' REPLY CLAIM CONSTRUCTION BRIEF**

## **TABLE OF CONTENTS**

Introduction .....	1
I. Disputed Terms in the '686 Patent.....	2
A. Modif[y/ied] [a] resource allocation / modify[ing] [the] computer resources allocated to a virtual server ('686 patent claims 5-7) .....	2
B. “resource unavailable messages” / “denied requests to modify a resource allocation” ('686 patent claims 5-7) .....	3
C. “determination that a virtual server is overloaded” ('686 patent claims 5-7).....	4
D. “virtual server” ('686 patent claims 5-7) .....	5
E. “determining that a second physical host can accommodate the requested modified resource allocation” ('686 patent claims 5-7) .....	7
II. Disputed Terms in the '726 Patent.....	9
A. “resource denials” ('726 patent claims 1, 4-5 & 8) .....	9
B. “quality of service guarantee” ('726 patent claims 1 & 4) .....	9
III. Alleged Means Plus Function Terms for the '726 and '686 Patents .....	10
A. '686 patent claim 7 “component” terms (i.e., clauses 1-3 of Ex. A) .....	10
B. '726 patent claim 1, 3, 4, 5 & 7 (i.e., clauses 4-8 of Ex. A) .....	12
The '752 Patent .....	12
IV. Disputed Terms in the '752 Patent.....	13
A. “exhausted” ('752 patent claims 1, 9 and 24) .....	13
B. “consumed” ('752 patent claims 1, 9 and 24).....	13
C. “service” ('752 patent claims 1, 3, 9 and 24).....	14
D. Means-Plus-Function Terms .....	14
V. Disputed Terms in the '051 Patent.....	15
A. “virtual server” (claims 1, 3, and 6).....	15
B. “physical interface[s]” (claims 1 and 3).....	16

C. “storing . . .” / “storing . . .” / “receiving . . .” / “determining . . .” / “determining . . .” / “sending . . .” / “using . . .” (claims 1 and 3) .....	17
D. “customer forwarding [table(s)/information]” (claims 1 and 3) .....	18
VI. Disputed Terms in the ’818 Patent.....	19
A. “hierarchical token bucket resource allocation” / “token(s)” (claims 1, 17, 30, 32 and 42) .....	19
B. “enforcing . . .” / “receiv[e/ing] . . .” / “classify[ing] . . .” / “compar[e/ing] . . .” / “forward[ing] . . .” / “buffer[ing] . . .” (claims 1, 17, 30, 32, 33, 37-39, 42).....	21
C. “maintain[ing] a connection over a network fabric” (claims 1, 17, 30, 32, 42) .....	22
D. “virtual [network/storage network] interface layer of an application server” (claims 1, 17, 32, 42) .....	23
E. Alleged Means Plus Function Elements (claim 17) (individually set forth in Ex. C) .....	24
Conclusion.....	25



## **TABLE OF AUTHORITIES**

### **Cases**

<i>Accent Packaging, Inc. v. Leggett &amp; Platt, Inc.</i> , 707 F.3d 1318 (Fed. Cir. 2013).....	20
<i>Graphon Corp. v. Autotrader.com, Inc.</i> , Case No. 2:05-cv-530 (TJW), 2007 WL 1870622 (E.D. Tex. June 28, 2007) .....	7
<i>Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.</i> , 381 F. 3d 1111 (Fed. Cir. 2004).....	2
<i>Netfuel, Inc. v. F5 Networks, Inc.</i> , 2017 WL 2834538 (N.D. Ill. June 29, 2017) .....	11
<i>Zeroclick, LLC v. Apple Inc.</i> , 2018 WL 2450496 (Fed. Cir. Aug. 20, 2019) .....	11

## Introduction

IV’s Opening and Responsive briefs demonstrate that its proposed constructions align with the use of the disputed terms in the context of the claims in light of the specification, and consistent with the prosecution history of each patent. That these constructions are correct becomes even more evident in the context of each disclosed invention, which sprang from companies at the cutting-edge of their technical fields. Three of the patents-in-suit, RE 44,686 (“the ’686 patent”), RE 42,726 (“the ’726 patent”) and RE 43,051 (“the ’051 patent”), were invented at Ensim Corporation, where the inventors were all highly experienced in the fields of cloud computing and virtualization. Dkt. No. 9 at par. 27, 31. General Magic, the original assignee of U.S. Patent No. 7,949,752 (“the ’752 patent”), was a pioneer in cloud computing. *Id.* at ¶ 20. 3Leaf Systems, Inc., where the inventions taught by U.S. Patent No. RE 44, 818 (“the ’818 patent”) were developed, was at the forefront of network virtualization. *Id.* at ¶ 35.

The extensive briefing has also revealed VMware’s strategic approach to claim construction. Rather than filter each term through established claim construction canons, VMware employs creative arguments in the hopes of fostering non-infringement positions. For instance, the term “virtual server” is at issue in the ’686 and ’726 patents as well as the ’051 patent. Despite the fact that the ’051 patent is unrelated to the other two patents (which are related), VMware asks the Court to consider evidence from both the ’051 patent and a patent incorporated by reference therein in construing the ’686 and ’726 patents. With respect to the ’752 patent, VMware asks the Court to wholesale disregard previous constructions of the very same terms made by an experienced Magistrate Judge in a prior Report and Recommendation. VMware takes the term of art “hierarchical token bucket” from the ’818 patent and argues that it should be limited to a specific, prior art algorithm that is not referenced by the patent’s written description. And on two occasions, VMware touts a citation as being from a paragraph and sentence without the word “embodiment”—while not disclosing that the preceding paragraphs explicitly characterize those cites as preferred embodiments. These techniques do not result in

proper claim constructions and should be rejected. Instead, the disputed terms should be given their proper scope as read in light of the intrinsic record.

**I. Disputed Terms in the '686 Patent**

**A. Modif[y/ied] [a] resource allocation / modify[ing] [the] computer resources allocated to a virtual server ('686 patent claims 5-7)**

IV's Proposed Construction	VMware's Proposed Construction
"modif[y/ied] set of functions and features of a physical host used in implementing tasks for the virtual server" / "modify[ing] a set of the functions and features of a physical host used in implementing tasks for the virtual server"	"modif[y/ied] [a] quality of service guarantee" / "modify[ing] [the] quality of service guarantee of a virtual server"

In its Responsive Brief, VMware confirms that its construction reads the term "quality of service guarantee" from the preamble of two claims in the '726 patent into these disputed terms for every asserted claim of both the '726 and '686 patents. *See* VMware Resp. Br. at 1, 12. In doing so, VMware asks the Court to ignore the inference that the patentee "intended his choice of different terms to reflect a differentiation in the meaning of those terms." *See Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.*, 381 F. 3d 1111, 1119 (Fed. Cir. 2004). VMware argues this inference should be overcome because the reissue application may have been prosecuted by an affiliate of IV and not the original applicants. To start, what entity prosecuted the patent has no bearing on claim construction. Moreover, the same PTO examiner evaluated the original application and the reissue application and agreed with the amendments made in the reissue patent, including those which replaced "quality of service guarantee" with the disputed terms.

This strange argument attempts to obfuscate the real issue at hand, namely, that VMware is asking the Court to significantly narrow the disputed terms by reading in a phrase that (a) appears in only select claims, (b) was amended out of most claims during prosecution and replaced with the disputed terms, (c) refers to a single exemplary embodiment, and (d) is a disputed term itself being separately proposed for construction. Accordingly, VMware's proposed construction should be rejected and, as detailed in IV's Responsive Brief Section 1(A), IV's proposal should be adopted.

VMware’s attacks on IV’s proposed construction are unavailing. VMware again misstates IV’s proposed construction claiming that IV “concedes” that it is asking the Court to construe ‘resource allocation’ as ‘resource.’ VMware Resp. Br. at 11. This misses the point. As more fully explained in IV’s Responsive Brief at 6, IV has not and is not asking the Court to read out ‘allocation’ from the disputed term but, instead, recognizes the term is accounted for because in order for any resource to be used by a virtual server that resource must necessarily already be allocated. Indeed, IV’s proposal would not have a different scope if it were to include the term: “modify an allocated set of functions and features . . . .”

**B. “resource unavailable messages” / “denied requests to modify a resource allocation” (’686 patent claims 5-7)**

IV’s Proposed Construction	VMware’s Proposed Construction
“an indication that a request by the virtual server cannot be immediately serviced” / “a request by the virtual server that cannot be immediately serviced”	“indications that requests by the virtual server for additional resources are either implicitly or explicitly denied, resulting from denied requests to modify a resource allocation”  see also construction of “modify a resource allocation”

After two briefs on these terms, VMware has yet to provide a single intrinsic record cite supporting its proposed constructions. VMware Br. at 7-8; Resp. Br. at 3-4. Instead, VMware continues to confusingly construe two constituent terms as one and incorporate the second term verbatim into its proposed construction. This is particularly important here because, as discussed at pages 7-9 of IV’s Responsive Brief, the claim requires these terms to have distinct meanings and be read in their proper sequence. VMware, however, changes the sequence of the claim elements and renders the latter half of the term superfluous. More specifically, rather than account for the fact that denied requests to modify a resource allocation are the things that generate resource unavailable messages which are in turn used to determine whether a virtual server is overloaded and that both collectively equate to ‘resource denials’ as described in the specification, VMware’s proposal ignores the second half of the disputed term.

IV’s proposed constructions on the other hand recognize that ‘resource unavailable messages’ and ‘denied requests to modify a resource allocation’ are related but distinct, and

attributes to each the appropriate scope based on the teachings of the specification. *See* IV Resp. Br. at 6-9. For example, the claims require that the disputed terms result in a determination that a virtual server is overloaded, which is described in the specification as being accomplished through monitoring resource denials. *Id.* at 8-9. The term ‘resource denials’ is further defined in the specification as “any request by the virtual server that cannot be immediately serviced.” *Id.* When compared to the claims of the related ’726 patent—which use the term ‘resource denials’ in place of the two disputed terms discussed here—it is evident that the term *denied requests to modify a resource allocation* equates to a request by the virtual server that cannot be immediately serviced, while the term *resource unavailable messages* equates to an indication of the same. *Id.* Thus, by examining the sequence of elements and the specification’s disclosures, it is apparent that rather than reading out portions of the disputed terms, IV’s proposal provides a complete and consistent definition that should be adopted.

**C. “determination that a virtual server is overloaded” (’686 patent claims 5-7)**

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	<p>“determination that an average number of resource denials for a virtual server is beyond a pre-configured threshold”</p> <p>See also construction of “resource denials”</p>

VMware claims that its proposal is drawn from “a clear definition of this term in the specification.” VMware Resp. Br. at 5. Knowing that it is improper to read limitations into the claim from preferred embodiments, VMware tries to bolster its sole supporting specification cite by stating that it does not use the word ‘embodiment’ in the sentence (or the paragraph) in which the definition is found.” *Id.* VMware, however, omits the fact that the paragraph it cites as support is immediately preceded by one that starts with “FIG. 2A is a flowchart of **an embodiment** of the overall process for dynamically modifying the resources of a virtual server.” *See* Ex. E at 5:29-30. VMware’s citation is obviously a continuation of the explanation of the preferred embodiment in Figure 2A as it moves through the flow chart depicted therein from step 210 to step 220. VMware’s argument to the contrary mischaracterizes the specification. In

short, the Court should not accept VMware's invitation to read a preferred embodiment into the claim.

VMware's criticisms of IV's well-supported position that this term should be given its plain and ordinary meaning are unpersuasive. As an example, VMware takes issue with IV's proposed construction of plain and ordinary meaning because "similar" terms like "virtual server overload signal" and "indication that a first physical host is overloaded" have been proposed for construction and IV is not arguing that those be construed as plain and ordinary. VMware Resp. Br. at 4-5. VMware does not cite—and IV is not aware of—any support for such a position. The disputed term is different than the other 'similar' terms that VMware cites, and unlike many similar terms, is understood by one of skill in the art as being used in its plain and ordinary meaning. Contrasted with VMware's construction based on one exemplary embodiment from the specification, IV's proposal is consistent with the intrinsic record. IV Br. at 18-20; IV Resp. Br. at 10-11.

**D. "virtual server" ('686 patent claims 5-7)**

IV's Proposed Construction	VMware's Proposed Construction
Plain and ordinary meaning, in the alternative "a virtual machine that resides on a physical server and uses the physical server's resources but has the appearance of being a separate dedicated machine"	"a process executing on a host computer that accepts communication requests, and that is capable of receiving a quality of service guarantee from a physical host"

Finding no support for its construction of this term as being "a process" in the '686 patent, the very first line of VMware's Responsive Brief focuses the Court on the use of "virtual server" in the '051 patent. *See* VMware Resp. Br. at 6. VMware then invokes a chain of alleged evidence that spans an unrelated patent, art incorporated by reference into that unrelated patent, and an examiner's characterization of prior art from the patent that the '686 patent reissued from. Even this artfully crafted chain of 'evidence,' however, cannot overcome the reality that VMware's proposed construction is entirely unsupported by the intrinsic record and contrary to the teachings of the '686 and '726 patents.

VMware’s characterization of a single citation as a “clear definition from the specification” is taken from the paragraph titled “DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS”—which makes no mention of the virtual server being a ‘process.’ *See* Ex. E at 3:46-58. Furthermore, far from being a “clear definition,” this isolated statement describes what a virtual server is “capable” of doing in the context of the specification’s preferred embodiments; it does not pretend to encompass the full scope of what a virtual server is. The two citations from VMware’s expert’s declaration used as support are also unpersuasive. The first refers the Court to the declaration’s section regarding the unrelated ’051 patent, and the second covers five pages of conclusory allegations as to the ’051 patent, supported by unrelated patents or inapplicable prior art. VMware Resp. Br. at Ex. 2, 13-14, 42-42.<sup>1</sup>

Similarly, VMware overreads the discussion surrounding the Yu reference during prosecution of the ’937 patent, which actually supports IV’s proposed construction. Yu disclosed a type of prior art virtual server that operated in a different manner and for a different purpose than the virtual servers in the ’686 and ’726 patents. *See* VMware Resp. Br. Ex. 46 at 1:20-29. VMware’s selected quotation is not the applicant characterizing his invention, rather, it is the examiner simply reciting the teachings of Yu. *See* VMware Resp. Br. Ex. 49.

VMware additionally attempts to overcomplicate the issue at hand by capitalizing on the fact that a virtual machine is software which is necessarily a flexible construct, and it is often difficult to conceptualize the point where software ends and something tangible like hardware begins. VMware uses this inherent flexibility to reduce the disputed term to the simplest example, a single process, and attempts to bolster that by pointing to prior art like Yu. But it cannot escape the fact that a process alone is merely a set of instructions being executed by a computer, something that by its very nature is not capable of performing all of the features of the claimed virtual server.

---

<sup>1</sup> These page numbers are VMware’s cites to the ECF page numbers.

Finally, VMware’s attempt to paint IV’s alternate construction as unsupported disregards the intrinsic record. While a plain and ordinary meaning construction is the most appropriate here, IV’s alternate proposal finds support throughout the specification starting with the title itself. The title of the patents—Dynamically Modifying the Resources of a Virtual Server—illustrates both the failings of VMware’s proposed construction and the merits of IV’s alternate proposal. A typical process alone is not capable of having resources. It uses resources while executing, but does not have ‘its own’ resources, particularly not dynamically modifiable resources. Furthermore, a typical process cannot perform many of the basic functions of the virtual server described in the Abstract, let alone the claims. A process cannot become overloaded based on resource denials, its resources cannot be dynamically modified to respond to changing resource requirements, it can’t even be ‘hosted’ on a physical server never mind transferred from one host to another if resources on the first host are insufficient. *See generally* Ex. R at 39-41; Ex. E at 2:48-3:20. These are all features of a fully virtualized machine, i.e. a virtual machine. It is undisputed that a virtual machine (1) resides on a physical server and uses that physical server’s resources, and (2) appears as a separate dedicated machine. Therefore, IV’s alternate proposal captures what is described as a virtual server in the ’686 and ’726 patents, whereas VMware’s construction primarily relies on its (faulty) arguments regarding the unrelated ’051 patent.

**E. “determining that a second physical host can accommodate the requested modified resource allocation” (’686 patent claims 5-7)**

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning; in the alternative: “determining that a second physical host can accommodate the requests by the virtual server that could not be immediately serviced”	Indefinite, <del>or in the alternative “determining that a second physical host can accommodate the denied request to modify a resource allocation”</del>

VMware alleges this term is indefinite due to a lack of antecedent basis. There is no such issue because the scope of the disputed term is readily ascertainable by one of skill in the art. *See Graphon Corp. v. Autotrader.com, Inc.*, Case No. 2:05-cv-530 (TJW), 2007 WL 1870622, at \*11 (E.D. Tex. June 28, 2007). *See also* Akl Decl. at ¶ 32. More specifically, VMware’s



indefiniteness theory hinges on the fact that “the requested modified resource allocation” term is singular while the antecedent basis for that term is plural, and thus, according to VMware, the claim offers “no guidance as to which one” of the denied requests to modify a resource allocation the disputed term is referencing. VMware Resp. Br. at 9. But reading the claim elements in their proper sequence demonstrates that while the preceding reference to “denied requests to modify a resource allocation” is plural, it is part of a larger element that results in a singular “indication that a first physical host is overloaded.” First the claim element in question requires ‘denied requests to modify a resource allocation’ which result in ‘one or more resource unavailable messages.’ The one or more messages are used to make a determination that the virtual server is overloaded, which in turn is used to determine whether a physical host is overloaded and send an ‘indication’ regarding the same. It is only after this indication is received that the system looks to determine whether another physical host can accommodate the overloaded virtual server. *See, e.g.,* Ex. E. at claim 5(a).

This reading of the claim language is consistent with the patent’s written description. One of ordinary skill in the art reading the claim and specification would see that resource denials are monitored by selective interception of system calls. *Id.* at 5:31-33. The interception of these calls allows the system to gauge the current resource usage with respect to a particular virtual server resource, including at what point the current allocation is fully used. *Id.* at 5:33-37. When a virtual server needs an amount of a particular resource, but has reached the maximum of its current allocation, a resource denial is generated. *Id.* Since the current allocation amount is known to the system, the amount requested in excess of that allocation is also known. *Id.* at 5:43-47. The specification describes various techniques to account for multiple denials, such as taking the average within a specific time window. *Id.* The amount above the current allocation is then used to check against the corresponding amount of that resource held unallocated by the virtual server’s physical host. *Id.* at 5:51-64. If the physical host does not have enough additional resources unallocated to satisfy the difference then an indication of a physical host overload is sent to the system and the system attempts to locate another physical host with

enough of the requested resource unallocated such that it could support the overloaded virtual server. *Id.* Accordingly, when determining a second physical host that can accommodate *the requested modified resource allocation* one of skill in the art would know that the antecedent basis for that term refers to the previous ‘requests’ as processed per the claim elements and reflected in the singular indication of a physical host overload, and also understand that the disputed term is being used according to its plain and ordinary meaning. *See* Akl Decl. at ¶ 32.

## II. Disputed Terms in the ’726 Patent<sup>2</sup>

### A. “resource denials” (’726 patent claims 1, 4-5 & 8)

IV’s Proposed Construction	VMware’s Proposed Construction
“indications that requests by the virtual server cannot be immediately serviced”	“indications that requests by the virtual server for additional resources are either implicitly or explicitly denied”

As it did with the ‘686 patent’s “determination that a virtual server is overloaded” term, VMware argues its sole specification support is “definitional.” VMware Resp. Br. at 14. And, again, VMware misleadingly tries to support that argument by saying “[n]either the sentence nor the paragraph in the specification from which VMware’s proposed construction is drawn uses the language ‘embodiment.’” *Id.* But that cite is clearly regarding a preferred embodiment as explicitly described in the immediately preceding paragraph. *See* Ex. F at 7:41-61 (“FIG 3 shows **an embodiment of one process** for determining whether an individual resource in a virtual server has reached its resource limit.”). The following paragraph from which VMware garners its support continues describing Figure 3 and even a cursory read shows it’s still discussing a single preferred embodiment. *Id.* Something that is “an embodiment” of “one process” can hardly be said to be a clear definitional statement.

### B. “quality of service guarantee” (’726 patent claims 1 & 4)

IV’s Proposed Construction	VMware’s Proposed Construction
“a guaranteed resource allotment which can be dynamically increased/modified”	“information that specifies a guaranteed amount of an assigned resource, and that can be dynamically increased/modified”

<sup>2</sup> The ’726 patent terms that overlap with ’686 terms discussed above can be found in Exhibit 7 to VMware’s Opening Brief. These claim terms should be construed as IV proposes for the reasons stated above with respect to the corresponding ’686 patent terms.

The dispute between the parties with respect to this term first centers on VMware's proposed "information that specifies" limitation. VMware's argument insisting that "quality of service guarantee" must be "information" conflates two distinct concepts. Because the 'quality of service guarantee' is implemented in a computer system with respect to a virtual server at least partially implemented via software it must at some level be represented by some type of information. The fact remains, however, that the disputed term is not a message, or indication as is the case with some of the other disputed terms, and therefore, should not be construed according to how a quality of service guarantee is represented, rather, the construction should describe what a quality of service guarantee is. See, e.g., Ex. F at 2:18-40. VMware's construction glosses over this important distinction.

The parties' second dispute with respect to this term relates to VMware's "assigned resource" limitation. VMware challenges IV's claim that the specification does not describe "quality of service guarantee" as "an assigned resource" and cites to column 2 lines 5 through 13 of the '726 patent which notes that a virtual server is "typically assigned a fixed level of resources . . . ." VMware Resp. Br. at 15-16. This portion of the specification, however, is describing the **prior art**, specifically, the short comings of the prior art that the applicant's invention is designed to overcome. Thus, the specification does not describe the invention's quality of service guarantee as being an assigned resource. This makes sense since the term assignment implies a fixed level of resources that is inconsistent with the invention's teaching of dynamic resources. In fact, this limitation also renders VMware's proposed construction internally inconsistent as it recognizes the dynamic nature of the invention. *Id.* at 15.

### **III. Alleged Means Plus Function Terms for the '726 and '686 Patents**

#### **A. '686 patent claim 7 "component" terms (i.e., clauses 1-3 of Ex. A)**

As discussed in IV's Opening Brief § III(B)(10)(i),<sup>3</sup> the structural disclosure of the first element and preamble include sufficient structure linked to the "component" terms that VMware alleges trigger the application of § 112 ¶ 6. *See Zeroclick, LLC v. Apple Inc.*, 2018 WL 2450496, \*3-4 (Fed. Cir. Aug. 20, 2019). VMware argues that there is no link between the structure IV points to in the claim and the claimed "component." VMware Resp. Br. at 10-11. This is not so. VMware fails to account for the fact that the term "component" is generally understood to mean "a constituent part." *See Netfuel, Inc. v. F5 Networks, Inc.*, 2017 WL 2834538, at \*5-8 (N.D. Ill. June 29, 2017). The structural elements preceding the component terms claim well-known structural components such as a 'virtual server operating in a first physical host of multiple physical hosts,' and at least one processor and memory. Therefore, when subsequently claiming a component configured to . . . that component is necessarily a constituent part of the previously disclosed structure. The previously disclosed structure here including, having a processor and memory and being situated in a system containing physical servers and virtual machines, both of which the component must interact with in order to perform the claimed function. This position is bolstered by looking to the language of each component element. For example, clause 1 of Ex. A claims that the component must receive an indication of a first physical host overload, i.e., must be communicatively coupled to the physical host server of the preamble. One of skill in the art would readily understand that as a whole this provides sufficient structure, and thus should not be subject to § 112 ¶ 6.

The remainder of VMware's arguments with respect to these disputed terms revolve around its position that the terms are subject to § 112 ¶ 6 and the corresponding structure disclosed as performing the claimed functions. IV has detailed its position on these issues in

---

<sup>3</sup> IV recognizes that it made inconsistent arguments in its Opening Brief regarding whether the preamble of this claim and claims 1, 4 and 5 of the '726 patent are limiting. This was not intended. In light of construing the term 'quality of service guarantee' and that the preambles' of these claims contain structure for certain elements, IV agrees that the preambles are limiting.

depth in its Opening Brief at 21-23 and Responsive Brief at 18-21, and thus will not re-hash them here as VMware's Responsive Brief simply reiterates its prior arguments.

**B. '726 patent claim 1, 3, 4, 5 & 7 (i.e., clauses 4-8 of Ex. A)<sup>4</sup>**

VMware broadly dismisses these terms as failing to recite sufficient structure to avoid means-plus-function treatment with generic statements regarding the terms being a "black box" and references back to its discussion of claim 7 of the '686 patent. Nothing new has been presented in its Responsive Brief and therefore, its arguments remain unpersuasive because they fail to take into consideration how these terms differ from those of the '686 patent as well as how one of skill in the art would understand the terms. IV lays out these deficiencies in detail in its Opening Brief at Section III(B)(10)(ii)-(vi) and its Responsive Brief at Section 3(B)(i)-(iv). In addition, should any of these elements be found to be means-plus-function, VMware's identification of the corresponding structure is too narrow and would impermissibly limit the claim scope as discussed in the aforementioned sections of IV's two prior briefs.

**The '752 Patent**

Because VMware asks this Court to construe the terms of the '752 patent with constructions already considered and rejected by Judge Mitchell, VMware begins its argument regarding this patent with a footnote downplaying the significance of Judge Mitchell's well-reasoned constructions based on the fact that the R&R had not yet been adopted by the District Court Judge and both parties filed objections to it. VMware Resp. Br. at n. 16. First, IV accurately described the procedural posture of the R&R to the Court's in its Opening Brief and noted that the R&R is persuasive, not determinative. IV Br. at 3. Second, that the R&R was not adopted or that it was objected to does not change the fact that a Magistrate Judge who has construed 537 claim terms during her time on the bench weighed the same evidence before this Court and disregarded each one of the arguments VMware re-hashes here.

---

<sup>4</sup> VMware lumps claim 4 in with claims 1, 3, 5 and 7, and then argues that the function of all the disputed claim elements from claims 1, 3, 4, 5 and 7 include "creating," a term only found in claim 4's "program code for" elements. This is improper for the reasons noted in IV's Responsive Brief at 22. Accordingly, IV addresses only claims 1, 3, 5 and 7 herein and refers the Court to its Responsive Brief for its position on the claim 4 elements.

#### IV. Disputed Terms in the '752 Patent

##### A. “exhausted” ('752 patent claims 1, 9 and 24)

IV's Proposed Construction	VMware's Proposed Construction
“used up to the allotted or pre-determined amount”	“unavailable for reuse”

The arguments made in VMware's Responsive Brief to support its limiting construction of this term were addressed in IV's Responsive Brief. IV Resp. Br. at 25-27. In summary, IV detailed that there (a) was no disclaimer made during the prosecution history, (b) is nothing in the specification equating the disputed term with ‘unavailable for reuse,’ and (c) is clear evidence in the specification showing that VMware's position is untenable by repeatedly disclosing that each agent is allowed to use up to a ‘pre-authorized’ amount of each service resource. *Id.*

Furthermore, VMware's most recent argument in its Responsive Brief that “claim 1 becomes incomprehensible when using IV's proposals for the ‘exhausted’ and ‘consumption’ terms” is contradicted in the very same sentence when it reads both proposals into the claim. VMware Resp. Br. at 25. Specifically, VMware claims that “used up to the allotted or pre-determined amount upon being used by the network-based agent” is “incomprehensible” because a resource would be “used up” upon being “used.” *Id.* In addition to this construction making perfect sense, it finds support in the patent specification. *See, e.g.*, Ex. D at 52:31-34, 22:34-37, 22:47-50, 24:63-67, 25:1-6. Since, as more fully described in IV's Opening and Responsive briefs, IV's construction reflects the intent of the patentee as disclosed by the intrinsic record, it should be adopted. *See* IV Br. at 4-5; IV Resp. Br. at 25-27.

##### B. “consumed” ('752 patent claims 1, 9 and 24)

IV's Proposed Construction	VMware's Proposed Construction
“used”	“used up”

VMware's Responsive Brief simply reiterates the flawed arguments it made in its Opening Brief regarding this disputed term. As detailed in IV's Opening and Responsive briefs, there is zero support for VMware's proposed construction and its attempts to sow confusion are easily dispatched. *See* IV Br. at 5; IV Resp. Br. at 28-29. In short, VMware has not and cannot

support its proposed construction because its arguments are belied by the prosecution history, the specification itself and the in-depth analysis of a United States Magistrate Judge. *Id.*

### C. “service” (’752 patent claims 1, 3, 9 and 24)

IV’s Proposed Construction	VMware’s Proposed Construction
“network functionality available to agents”	“an application that is used on behalf of a principal”

VMware’s arguments that the disputed term should be limited to “an application” because the ’752 patent’s exemplary services such as an “e-mail service” are software applications and not ‘network functionality’ should be disregarded. As described in IV’s Opening and Responsive briefs, such a position is untenable because it impermissibly narrows the claim based on a single embodiment<sup>5</sup> and relies on a twisted reading of ‘network functionality’ to try to discredit IV’s proposed construction. IV Br. at 6-7; IV Resp. Br. at 29. Furthermore, claiming that services such as ‘e-mail service[s]’ are not ‘network functionality’ is belied by the fact that e-mail, i.e. electronic mail, cannot function without a network so the services must necessarily be network functionality. *See, e.g.*, Ex. D at 5:32-5 (“[a] network system may be maintained, managed, and/or operated by any provider of technological services, such as electronic mail”) (emphasis added). Therefore, IV’s proposal accurately reflects the intrinsic record and gives the disputed term its full scope.

### D. Means-Plus-Function Terms

The parties’ primary dispute with respect to the ’752 patent’s means-plus-function claims stems from VMware’s inclusion of select citations from the specification into the agreed-to structural elements that correspond to the claimed function. As demonstrated in IV’s Opening and Responsive briefs, including these additional specification cites risks inappropriately limiting the claims. VMware’s Responsive Brief illustrates this flaw by noting that its original specification citation used to supplement the ‘service wrapper 26’ structural element read out at least one embodiment. VMware Resp. Br. at 29. As set forth in IV’s Opening and Responsive

<sup>5</sup> VMware also ignores the fact that software is not limited to an “application.” While an application is a specific compilation of software code designed to be executed to obtain a specific result, software itself is simply the opposite of hardware, e.g., a set of instructions for a computer.

briefs, VMware's additional citations are unnecessary and ultimately lead to more confusion and ambiguity, and therefore, should be disregarded. *See* IV Br. at 7-8; IV Resp. Br. at 29-30.

The only additional dispute involves the structure for the “means for monitoring an amount of the service resource used by the network-based agent.” IV has set forth its position on this issue in detail in both of its briefs. IV Br. at 8; IV Resp. Br. at 30-31. In short, VMware's structure fails to account for a portion of the claimed function, and its argument regarding the ‘converter’ not being part of the corresponding structure is based on a misreading of the record. IV Resp. Br. 30-31. For these reasons, the Court should adopt IV's proposal as it accurately reflects the scope of the corresponding structure.

## **V. Disputed Terms in the '051 Patent**

### **A. “virtual server” (claims 1, 3, and 6)**

IV's Proposed Construction	VMware's Proposed Construction
Plain and ordinary meaning  Alternatively – “virtual machine(s) that reside(s) on a physical server and use(s) the physical server's resources, but [has/have] the appearance of being a separate, dedicated machine(s)”	“a process executing on a host computer that accepts communications requests”

VMware's argument with respect to “virtual server” in the '051 patent reduces to limiting the term to a specific embodiment from the portion of the '051 specification that discusses prior art. Indeed, throughout its entire briefing on the term, VMware's two cites to the actual '051 patent specification can be found in a single discussion spanning only two short paragraphs of the Background of the Invention sections of the patent. VMware then uses this prior art section to bootstrap its arguments with respect to the patent application incorporated by reference in to the '051 patent (the '258 patent), and to commercial products used to implement a specific embodiment of the '258 patent. In doing so, VMware fails to capture the scope of the term as disclosed in the '051 patent and as understood by a POSITA at the time of invention.

VMware tries to connect the virtual servers disclosed in the '051 patent to the incorporated reference by stating that “[t]he '051 patent further explains how to create these private virtual servers by pointing to an incorporated reference that issued as U.S. Patent No.



6,976,258.” VMware Resp. Br. at 30. The ’051 patent does not make such a limiting statement. The citation referred to by VMware is found in the Background section of the specification and it states in non-limiting fashion that the referenced application discloses “a method for creating [] a private virtual server.” Ex. H at 3:64-67.<sup>6</sup> Not surprisingly, VMware cites to no support that would allow limiting this term to a single embodiment disclosed in the background of a patent. The remainder of the ’051 specification confirms that this construction is too limiting. For instance, the ’051 patent discloses “a host computer containing a plurality of virtual servers,” unlike the host computer disclosed in the ’258 patent that runs only a single “virtual host server.” *Id.* at 7:62-65. This is because the virtual servers described in the ’051 patent are more complex than a mere “process executing on a host computer that accepts communication requests.” For example, the ’051 specification discloses that “[e]ach private virtual server includes an associated IP stack 422. An IP stack is a set of software processes that together manage the transfer of information in packets according to Internet protocols.” Ex. H at 9:20-23. This is inconsistent with VMware’s contention that the inventors intended to limit the meaning of virtual server to a single, generic process executing on a physical computer.

As discussed in its Opening and Responsive briefs, a POSITA would readily understand the claimed virtual servers in accordance with their plain and ordinary meaning. IV Br. at 41-42; IV Resp. Br. at 31-33. Furthermore, IV’s alternate construction addresses deficiencies in VMware’s overly broad construction by clarifying that the functionality of a virtual server is abstracted from the underlying physical hardware but appears to be a fully functional physical server from the outside. *See, e.g.*, Ex. H at 4:42-45 (“[t]he privately addressed transmission may then be processed within the private virtual server in the same manner as if the server was actually a physical server resident within a LAN.”).

#### **B. “physical interface[s]” (claims 1 and 3)**

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	“hardware that provides a point of communication between two or more devices”

<sup>6</sup> VMware cites to column 4 rather than column 3. IV assumes this is a typographical error.

With respect to this term, VMware asserts that “IV is wrong that a physical interface is not a point of communication between two or more devices.” VMware Resp. Br. at 32. This line of reasoning, however, conflates a key distinction between physical interfaces themselves and the communication that takes place over them. A physical interface on its own does not provide communication, rather it provides physical connectivity to a physical medium over which communication can take place as implemented by, e.g., software. It would be misleading to insert a “point of communication” requirement into the definition of physical interface itself, because physical interfaces themselves do not necessarily provide communication without any additional programming or signals. This fact is bolstered by the remaining claim language, which includes “receiving a transmission” and “sending the transmission,” language that describes the communication functionality taking place over the underlying physical interface. *See* Ex. H at claim 1.

Regarding the proposed “between two or more devices” language, VMware argues that “there is no doubt that the physical interface is providing a point of communication between at least the remote location and the private network users.” VMware Resp. Br. at 33. However, VMware is still conflating the meaning of a physical interface with functions that a physical interface can support. For example, a physical interface is still a physical interface whether it’s connected to one, two, three, or zero devices, and may support functionality such as 802.11 scanning that occurs without being connected to any device at all. Again, IV is not suggesting that physical interfaces cannot be involved in establishing communication between devices, but rather that the definition of the term itself is well-understood and does not expressly include VMware’s proposed concepts of “two or more physical devices” or “provid[ing] a point of communication.”

**C. “storing . . .” / “storing . . .” / “receiving . . .” / “determining . . .” / “determining . . .” / “sending . . .” / “using . . .” (claims 1 and 3)**

Here, VMware has not proposed any claim term for construction but rather asks the Court to read limitations into each step of the asserted claims under the guise of providing “clarity.”

VMware Resp. Br. at 34. Not only is this request misplaced but it would unnecessarily limit and confuse the claims. As set forth in detail in IV's Opening Brief at 43-47 and Responsive Brief at 34-36, the claims, when properly read, do not suffer any antecedent basis problems and, therefore, do not require the insertion of the limitations VMware proposes. More specifically, the terms begin by generally describing "a" physical interface and "a" tunnel identifier followed by referring to those terms with the definite article "the." *See, e.g.*, IV Resp. Br. at 34. In addition, VMware's proposed insertions would limit the claims based only on exemplary embodiments from the specification, further highlighting that its construction should not be adopted. *Id.* at 36; IV Br. at 46-47.

VMware also continues to argue that its "proposal does not speak to whether a tunnel is uni- or bi-directional." This argument is facially inconsistent with its proposed construction, which would insert the directional terms "incoming" and "outgoing" to qualify every instance of physical interface and tunnel identifier recited in claims 1 and 3. As noted by IV throughout the briefing process, the specification contemplates transmissions and tunnels in either direction and/or in both directions. IV Resp. Br. at 36; IV Br. at 46.

**D. "customer forwarding [table(s)/information]" (claims 1 and 3)**

IV's Proposed Construction	VMware's Proposed Construction
"table(s) containing [a set/sets] of customer specific forwarding information" / "set(s) of customer specific forwarding information"	See construction of "storing a customer lookup table, the customer lookup table storing associations between physical interfaces and tunnel identifiers identifying tunnels for private networks and a plurality of customer forwarding tables" / "storing . . . customer forwarding information, . . . the customer forwarding information associating network addresses with physical interfaces and tunnel identifiers"

After not addressing this term in its Opening Brief, VMware states in its Responsive Brief that the term should have its plain and ordinary meaning. VMware Resp. Br. at 36-37. VMware takes issue with two aspects of IV's proposed constructions, which differentiate between singular and plural instances of the terms. *Id.* VMware's first argument is that IV's proposal is too broad because it does not provide a "reference to information associating network

addresses with outgoing physical interfaces and outgoing tunnel identifiers”—a limitation VMware argues is disclosed elsewhere in the claim. *Id.* at 36. If correct, the fact that the limitation is already present in the claim, however, is precisely why it should not be read into this disputed term.

VMware’s second criticism is that IV’s inclusion of “set(s)” of “specific” customer forwarding information is unnecessary because there is other language in the claims “reciting determination of ‘the correct’ customer forwarding table/information accessible to each customer.” *Id.* But this misses the point of IV’s proposed construction. *See* IV Resp. Br. at 37-38. IV’s proposal is made to simply clarify when the word “information” is used in its singular or plural sense since the word can have both meanings. VMware’s cite to the other claim language does not clarify whether information is being used in the singular or plural. And, as set forth more completely in IV’s Opening and Responsive Briefs, IV’s proposed construction is consistent with how a POSITA would read the claim in light of the specification. IV Opening Br. at 39-40; IV Resp. Br. at 37-38. VMware has not disputed the supporting intrinsic evidence relied on by IV with respect to this term.

## **VI. Disputed Terms in the ’818 Patent**

### **A. “hierarchical token bucket resource allocation” / “token(s)” (claims 1, 17, 30, 32 and 42)**

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	The specific class-based scheduling algorithm known in the art as the “hierarchical token bucket” / “token” as used in “hierarchical token bucket resource allocation”

Here, VMware’s Responsive Brief confirms that VMware’s proposed construction requires two steps, the second bootstrapping off of the first, and neither following proper claim construction procedure. VMware starts by remaking the generic term “hierarchical token bucket” into a proper noun (“just like the Statue of Liberty”) and then uses extrinsic sources to limit that proper noun to a specific algorithm, “the specific class-based scheduling mechanism known in the art as ‘the hierarchical token bucket.’” VMware Resp. Br. at 38-39. But, as set

forth more fully in IV's Opening and Responsive briefs, the patent does not limit hierarchical token bucket to a proper noun, let alone to a specific algorithm allegedly represented by the proper noun.<sup>7</sup> IV Br. at 31-32; IV Resp. Br. at 38-40. Instead, the patent utilizes the term to refer to the broader discipline or concept of a hierarchical token bucket resource allocation as understood in the art. IV Resp. Br. at 38-40. The claims themselves expressly recite a hierarchical token bucket to enforce the claimed bandwidth resource allocation scheme in a virtual I/O environment. Consistent with this requirement, the specification details the scope of the claimed hierarchical token bucket scheme in column 8, lines 20-39 by referring to the disclosure as how the term is taught "in the present invention."

Based on this disclosure and how the term is used in the claims, a POSITA would readily ascertain the intended meaning of the disputed term, i.e., the general discipline of hierarchical token bucket allocation. IV Resp. Br. at 38-39; Madisetti Decl. at ¶¶ 42-45. Further illustrating the point, the specification goes on to depict two exemplary hierarchical token bucket configurations in Figs. 12-13, neither of which match the specific algorithm known in the prior art that VMware attempts to read into the term.<sup>8</sup> This makes sense because it was with respect to these two figures where the patentee first tagged the disputed term as HTB, clearly indicating he was applying the concept—not a specific implementation—as illustrated in the aforementioned figures. VMware's construction would substitute a specific, extrinsic prior art algorithm for what is more broadly disclosed in the patent—a construction that is "rarely, if ever, correct." *Accent Packaging, Inc. v. Leggett & Platt, Inc.*, 707 F.3d 1318, 1326 (Fed. Cir. 2013).

---

<sup>7</sup> In fact, this 'proper noun' argument has no basis because VMware overlooks the fact that the patentee is not using HTB as a proper noun at all, but rather, for ease of reference tagged the general term 'hierarchical token bucket' as 'HTB' in column 2 line 61-62 in reference to figures 12 and 13. *See* '818 at 2:61-62 ("Figs. 12 and 13 are diagrams showing example hierarchical token bucket (HTB) hierarchy configurations.").

<sup>8</sup> *See also* Ex. G at 10:2-14; 11:51-52; 13:40-47 ("FIG. 12 illustrates one hierarchical configuration against which I/O communications can be classified" and "[a]s FIG. 13 illustrates . . . the hierarchical configuration is more granular . . . the leaf nodes, at the SAN level, correspond to virtual Host Bus Adapters associated with individual application servers, as opposed to storage and network classes."); 2:9-18 ("two-tier hierarchical QoS management process" employed in both ingress and egress directions); 9:51-67 (applying a hierarchical token bucket scheduling mechanism to implement exemplary QoS management process in context of claimed resource allocation scheme).

**B. “enforcing . . .” / “receiv[e/ing] . . .” / “classify[ing] . . .” / “compar[e/ing] . . .” / “forward[ing] . . .” / “buffer[ing] . . .” (claims 1, 17, 30, 32, 33, 37-39, 42)**

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	“enforcing . . . across the physical [storage network] interface of the virtual I/O server” / “receiv[e/ing] in the virtual I/O server” / “classify[ing] in the virtual I/O server” / “compar[e/ing] in the virtual I/O server” / “forward[ing] in the virtual I/O server” / “buffer[ing] in the virtual I/O server”

VMware argues that the limitation “. . . in the virtual I/O server” should be read into nearly every claim element of every independent claim at issue. In its Responsive Brief, VMware attempts to capitalize on two instances where the patent describes “the present invention” as teaching managing quality of service in virtual I/O servers. VMware Resp. Br. at 40. The fact that the patent utilizes the phrase “the present invention” is not significant for the reason cited by VMware since the patent does, in fact, claim a virtual I/O server where intended.<sup>9</sup> See, e.g., Ex. G at claims 32, 42 (expressly including language “in a virtual I/O server” in preamble only but not in claim body). Further, neither of the citations highlighted by VMware expressly describe each of the claimed elements as being performed “in the virtual I/O server” as advocated here by VMware. In its Opening Brief, VMware cites to descriptions of preferred embodiments, but VMware recognizes that even these descriptions do not support its attempt to read the term into virtually every element of the asserted claims. VMware Br. at 45-46 (“although the term ‘compare[e/ing]’ does not explicitly appear in Figures 5-10”). Thus, VMware is asking the Court to import a limitation into the claims based on an inference from a preferred embodiment.

<sup>9</sup> VMware’s reliance on *Honeywell* and *Edwards Lifesciences* is misplaced because both of those inventions involved a specific device (one a fuel filter and the other a medical graft) that were limited in application by the specific physical structure of the devices. Here, on the other hand, (a) the invention is not a **virtual I/O server** itself, and (b) because the invention is at least in part implemented in a virtualized computing environment, i.e., software, it is conducive to many embodiments, several of which are disclosed in the specification. The third case referenced by VMware is inapplicable to this issue.

As articulated in IV's previous briefs, each of the disputed claim elements at issue are unambiguous and explained in the specification. IV Br. at 33-34; IV Resp. Br. at 41-42. The addition of "in the virtual I/O server" does not clarify the underlying meaning of the claim language, rather it arbitrarily adds limitations contrary to any express requirements set forth in the specification. For example, the specification states that "[t]he present invention may be in the form of an I/O server application program . . . [a]lternatively, the present invention may be implemented at least in part by special purpose hardware such as ASICs." Ex. G at 3:53-58. The inventors thus contemplated several different structural configurations and variants for implementing the claimed invention including, but not limited to, a virtual I/O server, an application program, an ASIC, etc. It would thus be improper to substitute the "virtual I/O server" language in place of the express structural language selected by the inventors to express the claimed inventions.

**C. "maintain[ing] a connection over a network fabric" (claims 1, 17, 30, 32, 42)**

IV's Proposed Construction	VMware's Proposed Construction
Plain and ordinary meaning	"maintaining a connection between the physical interface of the application server and the physical interface of the virtual I/O server over a network fabric"

VMware argues that the term "maintain[ing] a connection over a network fabric" should be construed as "maintain[ing] a connection between the physical interface of the application server and the physical interface of the virtual I/O server over a network fabric." VMware's only support for inserting these limitations into the middle of the claim term consists of citations to preferred embodiments, which themselves would have to be strung together to reach VMware's proposed result. IV Resp. Br. 42-43. Faced with this, VMware asserts that its construction is the "only" "single meaning" consistent with the patent specification. But, the '818 patent uses open-ended language and terms of inclusion such as: "application server 102 according to one possible implementation of the invention," "FIG. 2 illustrates the protocol stack and modules of an application server 102 according to one possible implementation of the invention," "[i]n one



implementation, the virtual network interface layer 220 is configured with a link layer network interface profile,” etc. *See e.g.*, Ex. G at 3:41-42; 4:27-29; 5:21-24. Thus, there is no suggestion in the ’818 patent that the patentees intended to limit the claimed “connection” or “virtual interface layer” to the specific exemplary embodiments illustrated in the specification, and there are no additional words of exclusion that would otherwise require the disputed term to be construed as VMware proposes.

Putting aside the fact that VMware’s limiting proposal finds no intrinsic support, it fails for an even more basic reason, namely, it glosses over the fact that the invention is implemented in a virtual computing environment. Given this, a POSITA would readily understand that connections can be both physical and virtual and recognize that the inventors chose to define the structure of the claimed connection in terms of a virtual interface layer and not the physical interfaces that VMware contends. Madisetti Decl. at ¶¶ 48-49. This is made especially apparent considering that the patentees use “physical interface” elsewhere in the claim, illustrating that they did not intend to limit the disputed term as VMware suggests.

**D. “virtual [network/storage network] interface layer of an application server” (claims 1, 17, 32, 42)**

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	“a virtual storage network interface to higher layers of the virtual node in an application server” / “virtual network interface layer to higher layers of the virtual node in an application server” / “virtual interface layer to higher layers of the virtual node in an application server”

As an initial matter, VMware recasts IV’s citations, which demonstrated that VMware’s construction was too narrow, as IV proposing an alternative construction. VMware Resp. Br. at 44 (alleging IV argued that “virtual interface layer of an application server” means “interface layers [sic] that emulate layers of a networking or storage protocol stack.”). This is not correct. The citations that VMware points to simply illustrate that the disputed term is used throughout the specification as it is commonly understood by one of skill in the art. Despite VMware’s



attempt to redefine IV’s construction, the fact remains that IV has consistently maintained that the disputed term be construed according to its plain and ordinary meaning, and still does.

VMware now appears to have changed its proposed construction of the disputed term to “virtual interface layer (e.g., virtual network interface 220, virtual HBA 208a) of an application server that emulates a physical device or a virtual device relative to a physical device,” based on its twisted interpretation of IV’s opening arguments described above. Assuming this is the case, VMware’s new proposal fails for three reasons, (1) it simply takes a different exemplary embodiment and reads that into the disputed term impermissibly narrowing the claim, (2) it adds ambiguity to the claim because it is unclear what VMware means by “emulates a physical device or a virtual device relative to a physical device,” language that appears nowhere in the specification, and (3) as outlined in IV’s Opening and Responsive briefs, the disputed term has a clear and well-understood meaning to one of skill in the art and therefore further limiting it is inappropriate. Accordingly, VMware’s proposed construction—either the original or amended version—should be disregarded.

**E. Alleged Means Plus Function Elements (claim 17) (individually set forth in Ex. C)**

As discussed in IV’s previous briefs, claim 17 does not implicate § 112 ¶ 6. IV Br. at 36-38; IV Resp. Br. at 46-47. Claim 17 does not use any “means for” language that gives rise to the presumption of § 112 ¶ 6, and further fails to recite any prefatory nonce language that would overcome the presumption against means-plus-function claiming. VMware argues the claimed usage of “input/output virtualization modules” is equivalent to the unqualified “module for” language recited in *Williamson*. VMware Resp. Br. at 45-46. It further alleges that the holding in *ZeroClick* is irrelevant because the specification disclosed a prior art definition of “user interface code” whereas the ’818 patent allegedly fails to discuss prior art input/output virtualization. *Id.* However, VMware’s reasoning is flawed at least because the Background section of the ’818 patent describes that “[t]o reduce cost and increase flexibility for application servers to access available I/O subsystems, virtual I/O servers can be used to create logical

separations between the application servers and I/O subsystems to make the I/O subsystems as logical resource units to application servers.” Ex. G at 1:31-36. A POSITA would thus understand the I/O virtualization language as connoting structure associated with e.g., virtual I/O servers, application servers, and I/O subsystems as consistently described in the specification. *See* Madisetti Decl. at ¶¶ 50.

VMware’s assertion that “IV’s alternative identification of structure disclosed in the specification is unsupported” is conclusory and untrue. VMWare Resp. Br. at 47. IV’s identification of structure in the event that § 112 ¶ 6 applies is widely supported in the specification as articulated in its Opening Brief. IV Br. at 38. *See also* Ex. G at Figs. 5-10 (illustrating detailed, step-by-step walkthroughs of structured steps corresponding to various embodiments of the claimed invention). Accordingly, VMware has failed to show that claim 17 implicates § 112 ¶ 6 or, even if § 112 ¶ 6 is implicated, that the claim elements fail to recite sufficient structure as understood by a POSITA in the context of the specification.

### **Conclusion**

For the reasons stated herein and in IV’s Opening and Responsive Claim Construction Briefs, IV’s proposed constructions should be adopted.

Dated: April 10, 2020

/s/ Robert R. Gilman

DEREK GILLILAND  
State Bar No. 24007239  
SOREY, GILLILAND & HULL, LLP  
109 W. Tyler Street  
Longview, Texas 75601  
903.212.2822 (telephone)  
903.212.2864 (facsimile)  
derek@soreylaw.com

KARL RUPP  
State Bar No. 24035243  
NIX PATTERSON L.L.P.  
1845 Woodall Rodgers Fwy. Suite 1050  
Dallas, Texas 75201  
972.831.1188 (telephone)  
972.444.0716 (facsimile)  
krupp@nixlaw.com

**OF COUNSEL:**

Paul J. Hayes  
phayes@princelobel.com  
Matthew Vella  
mvella@princelobel.com  
Robert R. Gilman  
rgilman@princelobel.com  
Jonathan DeBlois  
jdeblois@princelobel.com  
Alex Breger  
abreger@princelobel.com  
PRINCE LOBEL TYE LLP  
One International Place, Suite 3700  
Boston, MA 02110  
Tel: (617) 456-8000  
Fax: (617) 456-8100

**COUNSEL for PLAINTIFFS**

**CERTIFICATE OF SERVICE**

I hereby certify that on the 10<sup>th</sup> day of April, 2020, I electronically filed the foregoing with the Clerk of using the CM/ECF system which will send notification of such filing to all counsel of record.

/s/ Robert R. Gilman